Norwegian School of Economics
Bergen, Spring 2017

Dynamic Complementarities in Infant Health

Ingrid Mikkelsen Semb

Supervisor: Aline Bütikofer

Master thesis, Economic Analysis

NORWEGIAN SCHOOL OF ECONOMICS

Acknowledgements

Working on my master thesis has been highly rewarding. I would like to thank my supervisor Aline Bütikofer for suggesting a very interesting topic, and for giving invaluable advice throughout the process. I am also grateful to Bütikofer and her co-authors Katrine Løken and Kjell Salvanes for generously making their data available for this thesis.

Bergen, June 2017

Ingrid Mikkelsen Semb

Abstract

A large literature documents that early-life shocks may substantially affect health and labor market outcomes. However, we do not know much about how different shocks interact. In theoretical work, the idea of dynamic complementarities has been put forward, implying that the effect of a positive shock is larger for individuals with a higher baseline level of skills. With dynamic complementarities, negative shocks decrease the returns to subsequent investments for affected children. In this setting, universal interventions may in reality serve to strengthen initial differences. The policy relevance of understanding whether dynamic complementarities exist can therefore hardly be underestimated, but no clear answer has yet emerged from the literature. I utilize two arguably exogenous sources of variation in infant health in Norway between 1935 and 1945 to shed light on this question. The first shock is fetal exposure to seasonal influenza, which has been shown to negatively affect long-term health and labor market outcomes. The second shock is access to well-child visits, which is known to improve similar outcomes. Specifically, I ask whether fetal exposure to influenza reduces the long-term returns to mother and child health care centers. This is a particularly relevant question for policy-making because seasonal influenza is a frequent shock, yet easily preventable by means of vaccination. If influenza lowers the returns to a universal intervention like mother and child health care centers, inequality of opportunity could therefore be reduced by increasing vaccination rates among pregnant women. Using detailed individual-level registry data, I find little evidence of influenza affecting the educational or labor market returns to well-child visits. This lack of significant interaction effects could however be driven by the two shocks not affecting the outcomes of interest in my sample. Further research is therefore necessary in order to understand whether returns to mother and child health care centers are lower for infants who were exposed to maternal influenza.

Contents

List of Figures

List of Tables

1 Introduction

A large and growing literature in economics focuses on how events in infancy and early childhood can have important impacts on adult health and labor market outcomes (see for instance Almond and Currie 2011; Almond et al. 2017). Understanding the effects of different early investments and their relative importance is key to reducing inequality of opportunity. However, while most of the previous literature has been aimed at evaluating single interventions, the policy importance of understanding how different programs interact can hardly be underestimated (Almond and Mazumder 2013).

Cunha and Heckman (2007) put forward the idea of dynamic complementarities, implying that skills produced at one stage of a child's life cycle raise the productivity of investment at subsequent stages, a feature often referred to as 'skills beget skills'. Whether dynamic complementarities exist is a question with important implications for policy design. Complementarities strengthen the case both for early investment, to enable satisfactory long-term outcomes in a cost-efficient manner, and for follow-up investments later in life (Cunha and Heckman 2007).

However, it is not clear from the literature whether dynamic complementarities exist. A rather frequent finding when assessing single interventions is that the worst-off benefit the most.[1] Although results from studies looking at heterogeneous effects of single interventions are not directly comparable to the literature on dynamic complementarities, they illustrate that skills may not always in practice beget skills.

This thesis contributes to the small but growing literature on dynamic complementarities by investigating whether there are complementarities in two infant health shocks. The

[1] For instance, Bütikofer et al. (2017) find that children from a low socioeconomic background benefited the most from access to well-child visits. Carneiro et al. (2015) find that a generous expansion of maternity leave benefits in Norway in 1977 had larger impacts for those children whose mothers would not have been able to take much leave without the reform. Galiani et al. (2005) find that the reductions in child mortality caused by privatizing water provision in Argentina in the 1990s were largest for the poorest municipalities.

first shock is in utero exposure to influenza, which has been shown to have long-run negative effects on education, earnings and adult health (Almond 2006; Schwandt 2017). The second health shock is access to well-child visits, which is known to have positive long-term effects on outcomes similar to those Almond and Schwandt consider (Bütikofer et al. 2017). Specifically, I ask the following question: Does fetal exposure to influenza reduce the long-term returns to mother and child health care centers?

To the best of my knowledge, this thesis is the first to look at whether there are long-term complementarities between two quasi-exogenous shocks to infant health in a developed country setting.[2] Seasonal influenza is a particularly relevant shock because of the frequency of the disease, and because the illness can be prevented through vaccination of pregnant women. If infants who are exposed to maternal influenza benefit less from universal interventions such as access to well-child visits, vaccination would be a simple way of reducing inequality of opportunity among infants. Although the setting and data are from Norway in the 1930s and 1940s, this study should therefore be of policy relevance today both in the context of well-child visit programs being rolled out in developing countries, and in the context of low vaccination rates in countries with universal well-child visit programs.[3]

The remainder of the paper is organized as follows: In Section 2, I present the conceptual framework of dynamic complementarities, influenza and well-child visits, and give an overview of the literature most closely related to my study. In Section 3, I describe my data. In Section 4, I explain the theory behind the empirical strategy I use to identify complementarities. I also present my main specifications. In Section 5, I present my results, followed by robustness checks. In Section 6, I discuss limitations to my analysis and provide suggestions for further research. Section 7 contains concluding remarks.

[2] Gunnsteinsson et al. (2016) look at short-term complementarities in infant health in a developing country. They combine a quasi-exogenous shock with variation from a randomized controlled trial.

[3] Although the World Health Organization (WHO) recommends influenza vaccination to pregnant women (WHO 2016), vaccinations rates among pregnant women remain low in many countries (see for instance Norwegian Institute of Public Health (NIPH) 2015).

2 Background

The literature on early development is huge, and the topics studied span from environmental health effects (e.g. Almond and Mazumder 2011; Bharadwaj et al. 2014; Bleakley 2007; Currie and Walker 2011) to the importance of quality care (e.g. Carneiro et al. 2015; Garces et al. 2002; Havnes and Mogstad 2011). For this reason, I focus only on the literature most closely related to my thesis. I start by formally defining the concept of dynamic complementarities and reviewing the existing empirical literature on the topic. Then, I explain how in utero exposure to influenza can be seen as a negative shock to the fetus, and what short- and long-term consequences we may expect from influenza exposure based on the literature. In the last part of this section, I describe the rollout of mother and child health care centers in Norway and why we may expect outcomes to improve for infants who gain access to well-child visits.

2.1 Dynamic Complementarities

The idea behind dynamic complementarities is that the productivity of an initial investment in a child's capabilities may increase with subsequent investments (Cunha and Heckman 2007). For instance, we may expect children with higher initial levels of skills to benefit more from access to preschools (Aizer and Cunha 2012).

Cunha and Heckman (2007) formally define dynamic complementarities as

$$\frac{\partial^2 f_t(h, \theta_t, I_t)}{\partial \theta_t \partial I_t'} > 0, \tag{1}$$

where h denotes the parental characteristics, θ_t the vector of skill stocks at time t, and I_t the (parental) investment in the child at time t.

With complementarities, how much the human capital accumulation process is affected by investments depends on the stock of skills. In this framework, importantly, universal interventions could serve to increase differences between children with different levels of skills.

For my analysis, I treat influenza as a negative shock to the stock of skills θ_t, and allow I_t to capture the well-child visits investment in the child. Apart from influenza being a negative shock, the interpretation remains the same as in Cunha and Heckman (2007). We observe dynamic complementarities if infants who never were exposed to influenza, and therefore presumably have a higher stock of skills, benefit more from access to well-child visits.

To causally identify complementarities, exogenous shocks to both the initial skill level (θ_t) and to the subsequent investment (I_t) are needed, as pointed out by Almond and Mazumder (2013). Since settings where 'the lightning strikes twice' are rare, not much is yet known about the empirical existence of dynamic complementarities. Below, I present an overview of a select number of working papers that use two convincingly exogenous shocks to estimate interaction effects.

For Danish infants born between 1930 and 1957, Rossin-Slater and Wüst (2016) study the long-term interaction effect of home visits during infancy and targeted high-quality preschool. To answer their research question, they use administrative data together with intertemporal and spatial variation in access to well-child visits and high-quality preschool childcare. They find that although both interventions had positive impacts on their own, access to well-child visits substantially lowered the value-added of high-quality preschool childcare. For years of schooling and male adult income, the positive effect of preschool was lowered by approximately 85 percent, pointing to high-quality preschool childcare and nurse home visiting programs being substitutes rather than complements.

For Romania, Malamud et al. (2016) utilize the legalization of abortion as an exogenous source of variation in parental investments. Presumably, cohorts born after the legalization should be subject to somewhat higher parental investments than cohorts born while abortion was still illegal. They combine the reform with exogenous variation in access to high-quality schools, stemming from slots being allocated strictly on the basis of transition scores. Although both shocks had positive impacts on individuals' scores on an important exam taken towards the end of high school, there is little evidence of positive interactions. The authors provide suggestive evidence that the lack of significant interaction effects may be due to behavioral responses by students and parents.

Johnson and Jackson (2017) consider two shocks to the childhood environment in the U.S. in the 1960s and 1970s. Their first shock is the Head Start program starting in 1964, which among other things aimed at improving preschool and health care access for poor children. The second shock is court-ordered school finance reforms starting in the 1970s, which reduced inequality in school spending across the U.S. Using individual-level data from the Panel Study of Income Dynamics, they find evidence of dynamic complementarities between the two shocks, with the Head Start program increasing returns to the school finance reforms.

Taking a different approach to measuring complementarities, Gilraine (2016) analyzes whether the effect of accountability schemes that hold schools accountable for their students' performance increases when children are exposed in multiple periods. Specifically, he analyzes the effect on test scores of the accountability schemes that followed the U.S. No Child Left Behind Act of 2001, and finds that being exposed to the program for a second year increases the effect of the first year exposure with 0.2-0.3 standard deviations.

A related strand of literature asks whether the damage caused by negative income shocks following natural disasters can be undone. Such mitigating effects of investments on negative shocks are inconsistent with the model of dynamic complementarities as defined

by Cunha and Heckman (2007), since the theoretical model posits that initial inequality increases with subsequent investments. This literature is therefore also highly relevant to this thesis.

Adhvaryu et al. (2016) find mitigating effects of conditional cash transfers on early disadvantage in Mexico around the year 2000. Specifically, the authors find that adverse rainfalls lowers the agricultural wage, so that children born in seasons with adverse rainfall have significantly worse educational and employment outcomes. However, randomized access to the conditional cash transfers from the Progresa experiment mitigates a large share of this damage. Each additional year of exposure to the cash transfers is estimated to reduce the negative impact of the adverse rainfall with 20 percent. For most specifications, this interaction effect is almost half the size of the coefficient on access to the cash transfers, pointing to substantial heterogeneity in the returns to the cash transfers.

Duque et al. (2016) study two similar shocks in a Colombian setting. Using administrative data for all cohorts born in the 1990s, they find that access to conditional cash transfers has a positive impact on children's educational outcomes, and that the effect is large enough to mitigate the damage caused by exposure to extreme weather during the first years of life. However, the effect of cash transfers does not seem to *depend* on whether the child had been affected by extreme weather.

Gunnsteinsson et al. (2016) look at potential complementarities in the health of infants during the first six months of life. The authors utilize the fact that in 2005, a tornado hit an area in Bangladesh that was already involved in a large-scale randomized controlled trial (RCT) on vitamin A supplements to infants. For male infants, they find that tornado exposure increased the frequency of severe fevers, and lowered nutritional status. However, access to the vitamin A supplement at birth completely protected infants from these harmful effects. Their evidence points to a mitigating role of the vitamin A shock on differences in infant health.

In sum, a consensus on whether complementarities exist has not yet emerged from the literature. Whereas some studies suggest that returns to investments are higher for those with higher levels of baseline skills, other suggest that returns are higher for those who have been subject to negative shocks. The literature therefore gives little guidance as to what to expect from the interaction between influenza and well-child visits.

2.2 How Does Influenza Affect an Individual's Stock of Skills?

Influenza is a contagious illness caused by influenza viruses (Centers for Disease Control and Prevention, National Center for Immunization and Respiratory Diseases (NCIRD) 2016). The viruses are believed to spread primarily through droplet transmission. Because the viruses are constantly mutating, immunity based on exposure to previous influenza infections is lost over time (NIPH 2016). During a typical winter season, five to ten percent of the Norwegian population will get influenza. Although the viruses may infect anyone, pregnancy makes women more susceptible to a severe disease course due to changes in the immune system, heart, and lungs (NCIRD 2017). For a more extensive presentation of seasonal influenza, I refer to Schwandt (2017).

Although the exact mechanisms by which influenza affects the fetus are not yet well understood, a large literature both in medicine and economics has established that fetal exposure to influenza infection may have severe short- and long-term impacts. Below, I describe a selected number of studies that, like this thesis, are concerned with how in utero exposure to influenza affects long-run outcomes.[4]

Almond (2006) broke ground by studying the long-term effects of the Spanish flu, an influenza pandemic that started in 1918. Using U.S. census data, he found that individ-

[4]For short-term effects, see for instance Bloom-Feshbach et al. (2011), who find that maternal influenza infections are associated with increased rates of miscarriages, or Currie and Schwandt (2013), who document a link between third trimester infections and shorter gestation length.

uals who were in utero during the outbreak had lower educational attainment, income and socioeconomic status than surrounding birth cohorts, in addition to increased rates of physical disability and transfer payments. Although Almond's study has been criticized for a potential selection bias due to fathers of affected cohorts on average having lower socioeconomic status than fathers of surrounding cohorts (see Brown and Thomas 2016), replication studies with data from countries largely unaffected by World War I have also found long-term negative effects of in utero influenza exposure (Lin and Liu 2014; Neelsen and Stratmann 2012; Nelson 2010).[5]

Studies of other pandemics have also found negative impacts, such as Kelly (2011) who used British data to analyze the impacts of the Asian flu in 1957 on outcomes for children up to the age of eleven. She found that in utero exposure to the pandemic negatively affected cognitive development as measured by test scores at the age of seven and eleven. For physical outcomes, she only finds a negative effect for children of mothers with certain characteristics. Birth weight is negatively affected if the mother either smoked before pregnancy or were shorter than 154 cm. Height at the age of seven and eleven is negatively affected if the mother smoked.

The negative impact of influenza exposure may persist across generations. Using Swedish registry data for children whose parents were affected by the Spanish flu, Richter and Robling (2016) show that the pandemic decreased educational attainment for the second generation, but only if the potentially exposed parent was the same gender as the child. Months of schooling decreased by 3-4 months for females and by 4-7 months for males, whereas the probability of attending college decreased by 3-5 percentage points (p.p.) for females and by 7-11 p.p. for males. The authors note that the effects seem to run both though biological mechanisms and through changes in the socioeconomic status of the first generation.

[5]Also in a Norwegian setting, pandemic flu has been shown to negatively affect long-term outcomes. A recent master thesis finds that fetal exposure to the Spanish flu reduced educational attainment for males, and that the effect was particularly strong in the poorest municipalities (Bakken and Husøy 2016).

All of the studies above are concerned with pandemic influenza. Pandemics are suitable natural experiments, both due to arguably being exogenous shocks, and due to the high infection rate. However, as pointed out by Schwandt (2017), because pandemics are caused by mutated viruses to which we are not immune, they are hardly preventable. In addition, pandemics are rare. Seasonal influenza, on the other hand, is a frequent illness that can easily be prevented by means of vaccination. Understanding the consequences of seasonal influenza may therefore be of large policy importance.

Schwandt (2017) was the first to show that also seasonal influenza has long-term negative effects on the fetus. Using administrative data for Danish individuals born between 1980 and 1993, he finds that infants who were exposed to influenza in utero earn nine percent less than their non-exposed counterparts. In addition, they are 3-4 p.p. less likely to participate in the labor market, and 35 percent more likely to be on welfare. Schwandt provides evidence that these labor market effects may be driven by lower cognitive skills among those affected by influenza. Even though birth weight and prematurity are also affected by influenza exposure, parts of the long-term harm are likely not to be observed at birth. Comparing results of two different approaches, Schwandt finds that cohort analysis yields similar results as tracking individual mothers who were hospitalized with influenza infections during pregnancy.

Both Schwandt (2017) and others have analyzed impacts of influenza by the trimester of exposure. Exposure in the third trimester is associated with shorter gestation and lower birth weight (Currie and Schwandt 2013; Schwandt 2017), whereas first trimester exposure has been linked to increased risk of miscarriages (Bloom-Feshbach et al. 2011; Schwandt 2017). Exposure during the second trimester, which is the trimester with strongest neural brain development, has been shown to affect education and labor market outcomes (Kelly 2011; Schwandt 2017). The second trimester is therefore the period of main interest in my analysis.

2.3 The Return to Investment in Well-Child Visits

There is a growing literature showing that access to well-child visits or other forms of infant health care has important consequences both in the short and long run. Moehling and Thomasson (2014) look at various public health interventions provided in the United States under the Shephard–Towner Act in the 1920s and find that nurse home visits was among the most effective interventions in reducing infant mortality. Chen et al. (2016) compare infant mortality rates in the United States and four European countries, and put forward nurse home visits as one intervention that could help bring down the relatively high rates of infant mortality in the United States.

A number of studies have used the nurse home visiting program in Denmark as a natural experiment to causally identify the consequences of access to well-child visits. Looking only at the short-run effects of the program, Wüst (2012) finds that access to home visiting nurses increased infant first-year survival rates with 0.5-0.8 percent. Furthermore, she finds evidence suggesting that the program reduced mortality from acute enteritis, pointing to promotion of breastfeeding and proper infant nutrition being important mechanisms behind the mortality reductions. More long-term, using administrative data, Hjort et al. (forthcoming) find improved adult health for the infants who were exposed to the program. Exposure to the program reduces the mortality between the age of 45 and 57. In particular, cardiovascular health is improved, with lower probability of being diagnosed with and dying from cardiovascular disease. The authors find evidence of heterogeneous treatment effects, with mortality reductions being larger for females, infants born in urban areas, and infants born in areas with worse baseline levels of infant health.

For Sweden in the 1930s, Bhalotra et al. (2016) find substantial short- and long-term mortality effects of a program that provided information, support and monitoring of infant care. In particular, the program provided information to mothers about sanitation and nutrition. The authors estimate that the program reduced the risk of dying in infancy with

24 percent, whereas the risk of death before the age of 75 is reduced with seven percent.

For Norway between 1936 and 1960, Bütikofer et al. (2017) show that access to infant care in the first year of life improved educational, labor market and health outcomes. The authors estimate that the completed years of schooling increased with 0.15 years, whereas earnings increased with two percent. Furthermore, exposed infants are taller as adults, and have fewer health risks at age 40, suggesting that improved nutrition during the first year of life could have been an important mechanism. The estimated effects are largest for children from a low socioeconomic background and for children from municipalities with poor health status. The program seems to have long-lasting consequences, with intergenerational persistence in education decreasing with ten percent following the program. For well-child visits, my study uses the same natural experiment as in Bütikofer et al. (2017).

In Norway, mother and child health care centers were rolled out during the first half of the 20th century following concerns about child health and infant mortality (Bütikofer et al. 2017). Starting with a few centers, the rollout intensified during the 1930s. By 1946, more than 400 centers had been established all over Norway (Schiøtz 2003).[6] These were easily accessible, and their services were free of charge. Although the centers were open to anyone, centers mainly targeted their services at poor families.

By providing both services aimed at the infant and services aimed at the parents, the centers targeted the infants' early-life environment from two sides at the same time (see Rossin-Slater and Wüst (2016), who make this point about a similar intervention in Denmark). Infants would be examined by doctors and nurses, and referred to a doctor or hospital if necessary (Bütikofer et al. 2017). Mothers would be provided with advice on infant nutrition. In particular, breastfeeding was promoted. Mothers were also taught how to make nutritious and safe milk formulas. Additionally, mothers were advised on how to

[6]See Bütikofer et al. (2017) for maps illustrating the spatial variation in access to centers.

reduce the risk of infant mortality, for instance through proper clothing and hygiene.

The uptake was rather large, with approximately 40 percent of all infants in treated municipalities being registered as receiving care at a mother and child health care center within the first year after an opening, and rates increasing to well above 60 percent within three years (Bütikofer et al. 2017). On average, a child was examined at the center three to four times during its first year of life. For a more detailed presentation of the mother and child health care centers and their background, see Bütikofer et al. (2017).

3 Data

This thesis combines a variety of data sources into pooled cross-sections of cohorts born between 1935 and 1945. Below, I describe the data in more detail.

3.1 Influenza

For information about influenza morbidity for the period 1935-1945, I use the Medical Statistical Reports stored in the Norwegian National Archives. The Medical Statistical Reports describe the health situation in each medical district per year, reporting among other cause-specific morbidity and mortality numbers. Since the data was not electronically available, I manually copied the monthly number of influenza cases per district to a spreadsheet.[7]

The district structure largely remained unchanged in this period, but some changes need to be taken into account.[8] These changes are handled as conservatively as possible, reducing

[7] As of Jan 01, 1935, Norway consisted of 377 medical districts (Statistics Norway 1937a). By comparison, there were 426 municipalities in Norway as of Feb 26, 2017, implying that on average, a medical district covered just a little bit more than one modern municipality.

[8] See Statistics Norway (1937a, 1939a,b, 1940, 1941a, 1943, 1946, 1948, 1949).

the variation in the data somewhat. Districts that merge between 1935 and 1945 are treated as one during the entire period, including the years prior to merging. Districts that split between 1935 and 1945 are reported together for the entire period. For a few districts, influenza numbers were available on a more detailed geographical level for all years. In these cases, I treat each area as a separate district. This approach results in 372 districts.[9]

The influenza numbers in the Medical Reports are not adjusted for population size. I therefore used historical population data to compute the number of influenza cases per 1,000 inhabitants. Since municipality-level population data was not available from Statistics Norway (SSB) for years prior to 1951, I used 1951 population data as point of departure, and then subtracted the net population change each year to arrive at yearly municipality-level population measures.[10] The exact details of this process can be found in Section A2. In order to minimize the measurement error from having only yearly population numbers, I assumed a linear population growth between January and December each year.[11] I then used historical overviews over the municipalities covered by each medical district to compute yearly population measures per medical district.[12]

Figure 1 presents how the nationwide influenza rate per 1,000 inhabitants developed between January 1935 and December 1945. The influenza rate shows strong seasonality, with clear spikes every winter. However, the magnitude of the spikes varies strongly. For instance, the infection rate was more than five times higher at the peak in 1939 than at the peak in 1938.

[9] For infant mortality data, I only use the years 1936-1938. Since fewer districts were affected by changes in this period, I create a separate influenza data set in order to keep as much of the variation as possible. The same approach then results in 376 districts.

[10] 1951 population data are retrieved from Statistics Norway (2017a), whereas the population changes prior to 1951 are retrieved from Statistics Norway (2015).

[11] Since the yearly population changes are rather small, the monthly population measures are not affected by whether I assume linear or exponential growth between January and December each year.

[12] Sources: Statistics Norway (1932, 1937b, 1941b, 1951).

Figure 1: Nationwide Influenza Prevalence per Month

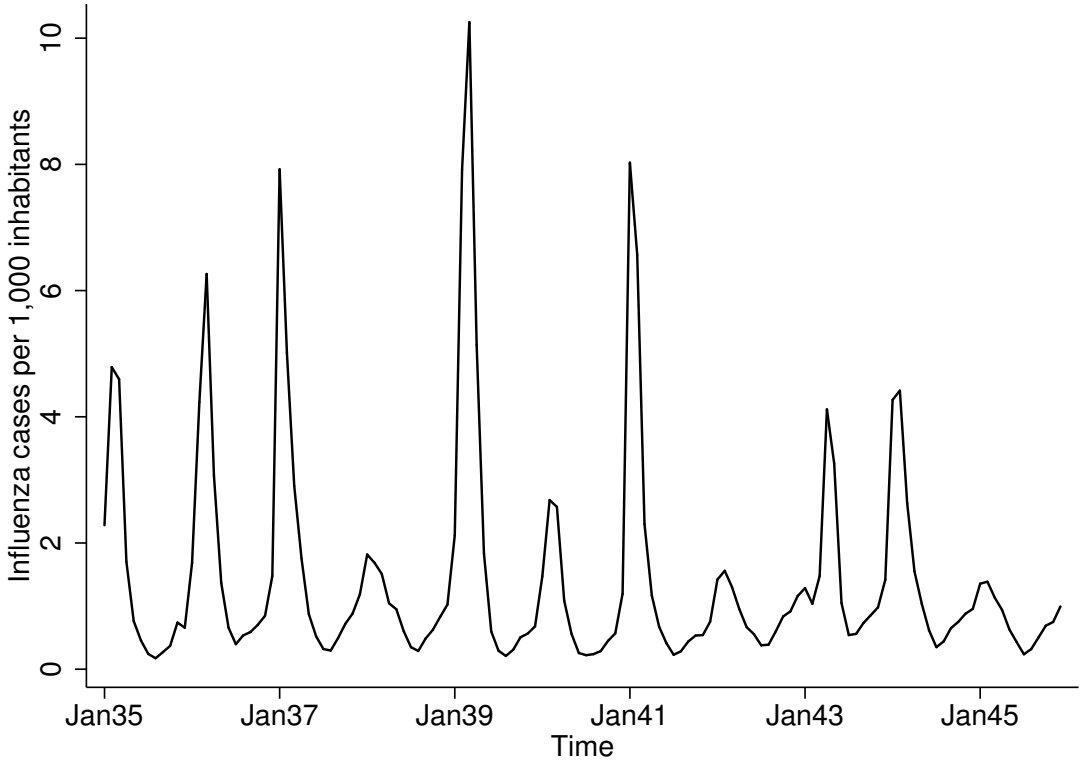

Notes: This figure shows the nationwide influenza prevalence per 1,000 inhabitants in Norway between 1935 and 1945, measured on a monthly basis.

It is also useful to look at geographical differences. Table A1 presents summary statistics for the influenza data for each county. The mean influenza rate is rather stable at 1-2 cases per 1,000 inhabitants for all counties. However, the maximum rate varies strongly across counties, with some counties having very large standard deviations. In Section 5.3, I therefore check how my results are affected by removing outliers.

The influenza data have important shortcomings, both in terms of the population data that is used to compute infection rates, and in terms of how the influenza was recorded. For the population data, the most important limitation is that it only accounts for how many people were formally registered in a given municipality. It therefore does not take the evacuation of the northern parts of Norway into account.[13] In order to check that my results are not driven by this limitation, I run alternative regressions on a subsample excluding the period affected by the evacuation.

The most important limitations of the influenza reports are that there may be systematic differences in how many of the actual cases are brought to a doctor's attention, which symptoms are recognized as influenza, and how accurately the doctors kept track of the number of cases. However, as I will explain in Section 4, district fixed effects should effectively deal with the part of the doctor-specific behavior that is constant over time. Another limitation is that I do not know which individuals were affected. I therefore need to assume that the overall influenza prevalence in a district is a good proxy for the share of infected pregnant women, meaning that the ratio of pregnant women to the overall number of reported cases should not vary substantially over time within a district.

[13] In the fall of 1944, the Nazi regime ordered large parts of the population of the northern counties Troms and Finnmark to evacuate (Petterson 2016). However, the approximately 67 000 inhabitants that moved south formally remained registered in their home municipalities. I have not been able to find any overviews over exactly where the evacuated population lived during this period, and the population numbers in the southern parts of Norway are therefore too low for the fall of 1944 and spring of 1945. I consider it likely that some municipalities sheltered a relatively larger share of the evacuated population than others, so that the measurement error may vary with unobserved variables.

3.2 Mother and Child Health Care Centers

The data describing the rollout of mother and child health care centers were collected by the authors of Bütikofer et al. (2017), and stem from surveys that were sent out to all Norwegian Women's Public Health Association (NKS) centers in 1939 and 1955. For a more detailed description of the collection process, see Bütikofer et al. (2017). These surveys contain information about the opening year of each center. For some centers, the opening month is also reported.[14] In my main specifications, I assume that centers opened in January if the opening month is not observable. I run alternative specifications using July to verify that assumptions on the missing opening months do not substantially change my results.

An important limitation of the mother and child health care center data is that it only covers centers run by NKS. Even though NKS was by far the most important center initiative in the period considered here, I cannot rule out that other initiatives established centers in the same period. Since the presence of other centers in municipalities not affected by NKS centers could bias my estimates towards zero, I follow the approach of Bütikofer et al. (2017) and include only the municipalities that got an NKS center before 1955 in my sample.[15]

3.3 Outcomes and Control Variables

For long-run data, I used Norwegian registry data from SSB. The registry data combines information from various official registries into a detailed individual-level data set covering the Norwegian population. In the following, I briefly describe the various variables.

[14] Bütikofer et al. (2017) do not consider the opening month in their analysis.

[15] Reassuringly, Bütikofer et al. (2017) provide evidence that the fraction of children registered as receiving care at a mother and child health care center in their municipality of birth goes from practically 0 to approximately 40 percent in the year of the first center opening by NKS. This suggests that the municipalities were not treated by other centers prior to the NKS center opening.

Information about gender as well as the municipality, year and month of birth are provided by the central population register.[16] Individuals born outside of Norway drop out of my sample because I use the medical district of birth in my identification strategy. Individuals born early in 1935 drop out because I do not have influenza data for 1934.

The central population register also provides mother and father identifiers, allowing me to control for whether the parents' education level and age at the time of the child's birth were above or below average. The mother identifier additionally allows me to include an indicator of whether the mother was unmarried in 1960.[17] Following Bütikofer et al. (2017), for all control variables, I keep the sample constant across specifications with and without control variables by including dummy variables indicating whether the control variable is missing.

The parental identifiers also serve to link siblings together. This is essential both in order to control for birth order, and to enable me to identify treatment effects by looking at children in the same family. When comparing siblings, having identifiers for both parents could potentially represent an improvement over only having a mother identifier because families where all children are born to the same father may be different than families involving multiple fathers. However, father identifiers introduce another measurement error, given that the tendency for mothers to report the correct father may depend on unobservable variables that also may affect outcomes of interest. Since the mother identifier is likely to be rather accurate, whereas the father identifier potentially has important limitations, I use only the mother identifier when comparing siblings.

Completed years of education is taken from SSB's educational database, or from the 1970 census for individuals who completed their education before the start of the education database. The census data is self-reported, but considered to be relatively accurate. Dis-

[16] The municipality of birth variable is based on the municipality structure around 1980. A detailed description of how these municipalities were traced back to the 1930s can be found in Section A2.

[17] Divorced and widowed mothers are not counted as unmarried.

counted average yearly earnings between 1967 and 2010 stem from the tax and income registry, and is given in 1998 NOK. The earnings measure is based on labor earnings, taxable sick benefits, unemployment benefits, parental leave payments and pensions. It is not top-coded. In order to enter in the income data, individuals need to be alive in 1967, and to earn taxable income between 1967 and 2010. The population covered by the registry data is therefore likely to be slightly positively selected.

Although registry data is likely to be much more reliable than self-reported data, there is some potential for measurement error. The income measures are based on reported taxable income only, potentially making it subject to under-reporting. The degree of under-reporting may depend on unobserved variables that may in turn affect the outcomes of interest, such as preference for risk.

For every district, I have created average measures by weighting the municipality level characteristic by the municipality's relative size in the medical district. The ratio of doctors to inhabitants in the child's year of birth was collected from SSB's historical yearly health statistics. The ratio of students to teachers in the child's year of school enrollment was collected from SSB's historical yearly school statistics.

Summary statistics are presented in Table 1. Average earnings for males is approximately 207,100 NOK, whereas the female average is less than half of that amount. Average years of schooling is rather high, with a large share of individuals taking more than the minimum level of schooling. The gender difference is much lower for education.

In order to look at potential short-term selection effects, I use data on stillbirths as well as mortality during the first year of life. Statistics Norway (2015) contains all reported live- and stillbirths in the period, allowing me to compute the stillbirth rate per municipality per year. A clear limitation of these data is that years with less than three stillbirths in a municipality are listed as missing for privacy reasons. Bakken and Husøy (2016) point to other reasons why the number of stillbirths is an imperfect measure. Examples include

Table 1: Summary Statistics

Outcomes	**Males**	**Females**
Earnings 1967-2010	204,044	88,293
	(111,522)	(66,433)
Observations	81,003	79,082
Earnings ages 31-40	230,518	71,369
	(88,563)	(69,359)
Observations	75,449	73,505
Earnings ages 41-50	251,868	119,331
	(127,300)	(82,531)
Observations	81,003	79,082
Years of schooling	11.256	10.313
	(3.151)	(2.693)
Observations	60,486	63,932
More than seven years of schooling	0.862	0.861
District background characteristics	**Both genders**	
Students per teacher	21.019	
	(7.418)	
Doctors per inhabitants	0.046	
	(0.102)	
Individual background characteristics		
Mother's age below mean	0.555	
Father's age below mean	0.535	
Mother's education below mean	0.743	
Father's education below mean	0.641	
Mother unmarried in 1960	0.195	

Notes: This table presents summary statistics for long-run outcomes and background characteristics. The variable 'Years of schooling' describes the number of completed years of education per individual. The variable 'More than seven years of schooling' indicates whether individuals took more than the seven years of schooling that were mandatory at the time. The variable 'Earnings all years' describes the average discounted yearly earnings between 1967 and 2010 in 1998 NOK. The variables 'Earnings 31-40' and 'Earnings 41-50' describe the average discounted yearly earnings between the age of 31 and 40, and 41 and 50, respectively. The variable 'Students per teacher' describes the ratio of students to teachers in the child's year of school enrollment, whereas the variable 'Doctors per inhabitants' describes the ratio of doctors to inhabitants in the child's year of birth. Both are measured at the medical district level. The individual background characteristics are dummies indicating whether the characteristic of interest is lower than the sample mean, except from the variable indicating whether the mother was unmarried in 1960.

differences across doctors in how long after the birth the infant has to remain alive in order to be counted as live-born, or what is considered a miscarriage and what is considered a stillbirth. Even though such differences should be less of an issue when including district fixed effects, I interpret estimates based on the stillbirth data with caution.

In order to compute infant mortality rates, I manually copied yearly district-level mortality numbers concerning children below the age of one from the Medical Statistical Reports to a spreadsheet. Complete data was collected for 1937 and 1938. I then used data from Statistics Norway (2015) on the number of live births per municipality in 1937 and 1938 to compute the number of deaths per 1,000 infants.

4 Empirical strategy

4.1 Theoretical Framework

This thesis aims to identify the causal interaction effect between influenza and mother and child health care centers. The inherent challenge with causal analysis is that each individual is observable only as treated (Y_{1it}) *or* as untreated (Y_{0it}), so that the individual causal effect of treatment, formally defined as $Y_{1it} - Y_{0it}$, is never observable (Angrist and Pischke 2009). Instead, we turn to studying outcome differences between treated and non-treated individuals. However, unless the outcomes of the two groups would have been equal absent the treatment, outcome differences cannot solely be attributed to the treatment. Differences in the composition of the two groups may lead to what we call a selection bias, which formally can be defined as differences in the expected outcomes absent treatment.

To identify the interaction effect between maternal influenza exposure and access to mother and child health care centers, we therefore want the treatment group and the

control group to differ only in terms of their influenza and center status. In an ideal experiment, we would randomly assign both influenza rates and mother and child health care centers across Norway. With large-scale randomization, we would expect the groups to be statistically indistinguishable, so that any differences in outcomes would be due to the treatment.[18] This is neither practically nor ethically possible. However, the rollout of mother and child health care centers in Norway around 1930, combined with variation in the influenza rate, provides a natural experiment that goes a long way in removing any selection bias.

Below, I will explain how we can use econometric techniques to create suitable control groups that allow us to view influenza and mother and child health care centers as exogenous shocks. First, I will explain how the effect of the centers can be analyzed with a rollout strategy. Then, I will explain how district, cohort, and mother fixed effects can be used to control for unobserved differences between districts and mothers with different levels of influenza.

4.1.1 Rollout

Since the rollout analysis is a generalized difference-in-difference (DiD) approach, I start by explaining the theory behind the DiD framework. For simplicity, I use the simplest possible setup: One treatment group ($D = 1$) and one control group ($D = 0$), and two periods: Pre-treatment, $t = 0$, and post-treatment, $t = 1$.

A DiD strategy assumes that without treatment, the outcomes of the treatment and the control group would have moved in parallel. In Figure 2, this assumption is graphically represented as a constant difference between the outcome of the control group ($E[Y|D =$

[18] Notably, Deaton (2010) argues that randomization should not be thought of as a 'gold standard'. It is beyond the scope of this thesis to discuss whether observational or experimental data would theoretically have been preferable. I therefore simply note that my research question could not be studied with experimental data.

Figure 2: A Theoretical Difference-in-Difference Model

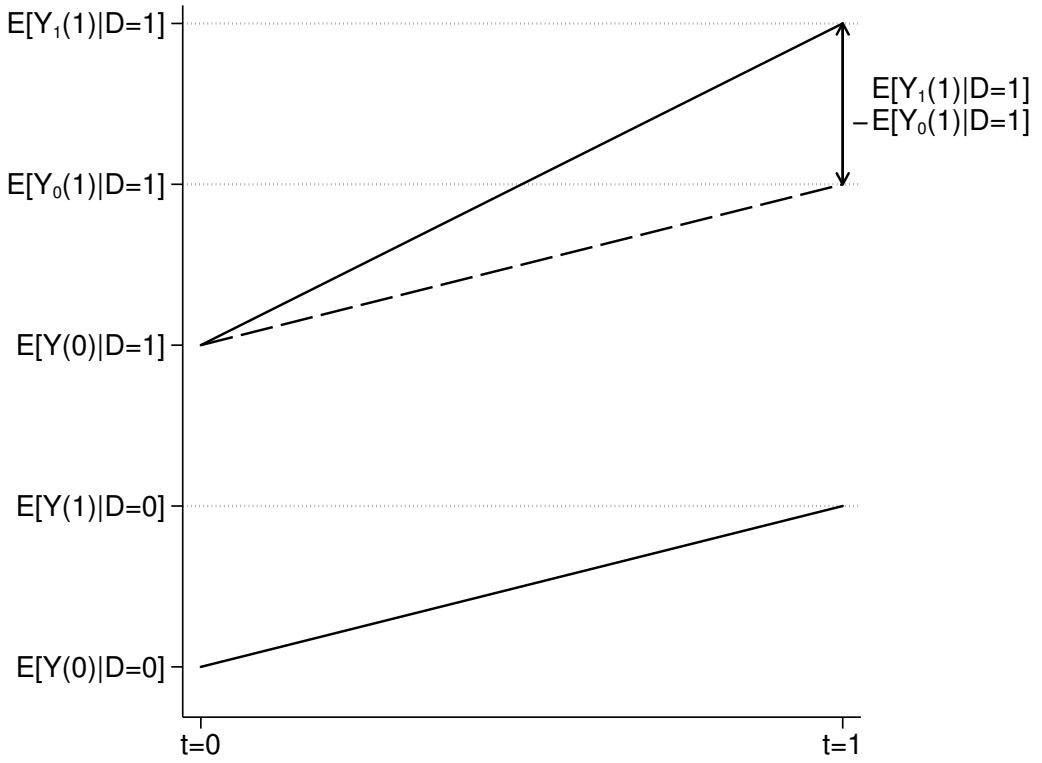

Notes: This figure illustrates a simple DiD setup with two periods and two groups. t indicates whether we are in the period prior to ($t = 0$) or post ($t = 1$) treatment. D indicates whether the group was treated ($D = 1$) or not ($D = 0$). The solid lines indicate the actual outcomes. The key assumption of the DiD framework is that without any treatment, the outcome of the control group would have moved in parallel with that of the control group, as indicated by the dashed line. Under this assumption, the vertical arrow indicates the average causal effect of treatment on the treated.

$0]$) and the counterfactual outcome of the treatment group ($E[Y_0|D = 1]$).

The identifying assumption of common trends is clearly a strong one. The treatment and the control group may be differentially affected by factors other than the treatment, such as an economic downturn. There may also be local changes affecting either the treatment or the control group, such as improved housing quarters for factory workers. It is therefore important to ensure that the treatment and the control group are as similar as possible, and to control for remaining differences to the extent possible. Time-invariant characteristics will be captured by the district fixed effects, as I will explain below. Controlling for char-

acteristics that vary over time introduces a potential problem since such characteristics may be affected by the treatment, introducing endogeneity.[19] I therefore use the common solution of interacting time-invariant characteristics with the time indicator, allowing me to control for differential linear trends without introducing endogeneity.

Under the assumption that outcomes would have moved in parallel absent any treatment, the counterfactual outcome of the treatment group is given by the outcome of the control group plus any initial difference between the groups. When the counterfactual outcome is known, the treatment effect indicated by the vertical arrow in Figure 2 is easily computed. Formally, the treatment effect[20] is given by:

$$\begin{aligned}\alpha_{TOT} &= E[Y_1(1) - Y_0(1)|D = 1] \\ &= (E[Y(1)|D = 1] - E[Y(1)|D = 0]) - (E[Y(0)|D = 1] - E[Y(0)|D = 0)\end{aligned}$$

For the mother and child health care centers, we need a slightly more generalized approach than the simplified example above. As described in Section 3.2, we only want to include districts that eventually got a center, meaning there is no suitable control group that never receives treatment. However, we can exploit the fact that center openings happened over time to create a control group based on the districts that have not yet gotten a center, but who will do so eventually. With a rollout strategy, districts are defined as part of the control group until they get a mother and child health care center.

[19] I talk more about the problem of 'bad controls' in Section 6.

[20] In this simplified example, since everyone assigned to the treatment group takes up the treatment, the treatment effect we estimate is a so-called treatment on the treated (TOT), meaning it describes the average treatment effect among those taking up the treatment (Angrist and Pischke 2009). In real-world examples, cases with a take-up rate of 100 per cent are rare. The identified effect is then a so-called intention to treat (ITT) estimate, which is the average effect among everyone in the treatment group, regardless of whether or not they actually took up treatment. With low take-up rates, we necessarily overestimate the extent to which the treatment group was treated. In such cases, ITT estimates underestimate the effect on those actually treated, and we divide the ITT estimate by the take-up rate to get the TOT estimate. In my case, since I do not know who were infected with influenza or who visited the mother and child health care centers, my estimates will be ITT in nature, and therefore constitute a lower bound of the TOT effect.

The crucial assumption behind a rollout is that the timing of the rollout does not depend on the potential for treatment, or on other factors affecting the outcome of interest. Bütikofer et al. (2017) perform a number of checks to verify that this key identifying assumption holds up for the Norwegian mother and child health care center rollout. First, they go through all relevant reforms and law changes in the period to check that they are not correlated with the timing of the center openings. They find no reforms or law changes that could potentially have been correlated with center openings for the birth cohorts I study. Secondly, they test whether the timing of center openings could be predicted by either 1930 district characteristics or changes in district characteristics between 1930 and 1946. The only characteristics they found to have predictive power were the 1930 urbanity status, population size, and inhabitants per doctor. They therefore exclude the largest cities at the time from their main sample, and control for inhabitants per doctor in the district at the time of birth. The same approach is followed in this thesis. Lastly, they use an event-study specification to check whether the timing of center openings depend on pre-opening trends. They find practically no anticipatory effects, pointing to the assumption that timing of the rollout does not depend on other factors affecting the outcome of interest holding up. Using a rollout strategy, I should therefore correctly identify the effect of getting access to mother and child health care centers.

4.1.2 District, Cohort and Mother Fixed effects

For influenza, we may worry that the infection rate is not a completely exogenous shock. For instance, we may expect crowded housing quarters to allow the influenza virus to spread more easily. Crowded quarters may also affect the outcomes of interest for instance through slower accumulation of human capital because children have nowhere to do their homework. If we fail to control for such variables that both affect the outcome and are correlated with regressors, the error term no longer has a zero conditional mean, and the OLS estimator is no longer unbiased (Wooldridge 2012). In our case, failing to control

for the fraction of the population living in crowded quarters could lead us to overstate the detrimental effect of influenza on long-term outcomes.

However, we do not have data on housing conditions. Instead, we may use fixed effects to control for unobserved characteristics that are constant over time. To see this, we may imagine a simplified setup where influenza is as good as randomly assigned across medical districts conditional on the unobservable fraction of people living in crowded quarters. Absent influenza, we therefore expect outcomes to be the same for those who get and those who do not get influenza, conditional on the housing conditions. Formally, $E(Y_{0dt}|H_d, D_{dt}) = E(Y_{0dt}|H_d)$ where the subscript 0 indicates the outcome without influenza, d indicates the district and t the time, whereas D_{dt} denotes the influenza status and H_d denotes the housing conditions.

Since housing conditions enter without a time subscript in the equation, and therefore by assumption is constant over time, we can implicitly control for it by including a district fixed effect in the regression model. District fixed effects capture all district characteristics that are constant over time, so that the relationship between influenza and the outcome of interest is estimated using only districts where the influenza infection rate varies over time.

In a regression setup, it is easy to see that lumping the unobserved H_d together with the zero-mean error term u_{dt} would give us an error term $\epsilon_{dt} = u_{dt} + H_d$ that would violate the zero-conditional mean assumption. Since we cannot simply move H_d out of the error term by directly controlling for housing conditions, we instead include the fixed effect $\alpha_d = \alpha + \gamma H_d$ and estimate the following model:

$$Y_{dt} = \alpha_d + \theta D_{dt} + u_{dt}$$

Now, under the assumption that the causal effect of influenza exposure on the outcome of interest is additive and constant, θ has a causal interpretation.

We can use fixed effects for units other than districts as well. In general, fixed effects control for unobserved characteristics that vary either with the time dimension or across the units we observe, and are constant along the other dimension. For instance, year fixed effects control for characteristics that affect entire cohorts in a given year, such as schooling reforms. With year fixed effects, the relationship of interest is estimated using only variation across units within each year. Family fixed effects control for effects that are constant across siblings, for instance the parents' preference for schooling. The relationship of interest is then estimated based on families where at least one child had access to the treatment, and at least one child did not.

There are different ways to include fixed effects in a model. One way is to include dummies. In the examples above, this would imply including dummies for all districts, years or families, respectively. Alternatively, fixed effects can be estimated as deviations from the mean, so-called within-group estimation. In the influenza example, we would calculate district averages and subtract them from the model:

$$Y_{dt} - \bar{Y}_{dt} = (\alpha_d - \bar{\alpha}_d) + \theta(D_{dt} - \bar{D}_{dt}) + (u_{dt} - \bar{u}_{dt})$$

Since α_d is constant over time, it falls away. The last possibility is differencing:

$$\Delta Y_{dt} = \Delta\alpha_d + \theta\Delta D_{dt} + \Delta u_{dt}$$

Again, since α_d has no time dimension, it falls away.

4.2 Identification Strategy

In order to analyze how influenza interacts with access to mother and child health care centers, I combine the empirical approaches from Bütikofer et al. (2017) and Schwandt (2017), with some modifications. Infants are defined as treated by mother and child health

care centers if they are born within 12 months of, or after, the first center opening in their medical district. Influenza treatment is defined as how much the influenza prevalence in the medical district of birth deviates from the typical influenza seasonality pattern in a period of interest during the pregnancy, since the deviation is likely to be less predictable than the actual infection rate (Schwandt 2017).[21] Whereas Schwandt uses month fixed effects to estimate the nationwide seasonality pattern, I have data on a more detailed geographical level and therefore include month-district fixed effects to estimate seasonality patterns per district. Following Schwandt (2017), when analyzing effects by trimester of exposure, I include all trimesters.[22] Since I cannot observe the exact period of infection, and since infection rates in different months are likely to be strongly correlated, including only the trimester of interest could potentially lead to an omitted variable bias.

The risk of influenza and of severe consequences may not be constant across mothers. Children born to different mothers may therefore not be directly comparable. However, Schwandt (2017) provides suggestive evidence that including mother fixed effects effectively controls for such maternal selection into influenza.[23] I therefore include mother fixed effects.[24]

[21] On the downside, this approach reduces the amount of variation in my data, and potentially increases the measurement error since small variations in the data are given more weight (Schwandt 2017).

[22] Since I do not know the date of conception, measures of influenza exposure during the different periods of pregnancy are calculated by counting months backwards from the date of birth of each infant.

[23] There is also another reason why Schwandt (2017) controls for mother fixed effects. In modern settings, mothers giving birth in different months have been shown to differ on important characteristics such as socioeconomic status (Currie and Schwandt 2013). However, as pointed out by Bakken and Husøy (2016), reliable birth control methods were not available in the period I study. Controlling for month of conception should therefore not be necessary in my setting.

[24] Whereas mother fixed effects are important for the internal validity of my analysis, it is important to bear in mind that effects estimated based on sibling samples may not necessarily carry over to the population of only children or twins.

Specifically, I estimate the following model:

$$Y_{idt} = \alpha_0 + \alpha_1 Flu_{dt} + \alpha_2 Center_{dt} + \alpha_3 (Flu_{dt} \times Center_{dt}) + \beta X_{idt} + \lambda_d + \theta_t + \gamma_{md} + \eta_f + \rho_d t + u_{idt} \quad (2)$$

where Y_{idt} is the outcome of interest for individual i born in district d at time t, and Flu_{dt} describes the influenza rate in district d in a given period of the pregnancy. $Center_{dt}$ is a dummy indicator equal to 1 if an NKS center opened in district d within a year after the individual was born, and 0 otherwise. X_{idt} is a set of individual and district-level characteristics. λ_d are district fixed effects, capturing time-invariant characteristics at the district level. Similarly, θ_t are cohort fixed effects, capturing common time shocks. γ_{md} is a set of district-month fixed effects, capturing the typical district-level pattern of influenza seasonality. η_f is a set of family fixed effects. ρ_d is the coefficient of a district-specific time trend variable t, which controls for districts potentially being on different trends. u_{idt} is the zero-mean error term.

α_1 measures the effect of influenza in a period of interest during pregnancy, conditional on not having a mother and child health care center in the medical district. α_2 measures the effect of having access to a center, conditional on no influenza. The coefficient of interest, α_3, measures how the effect of having access to a center is affected by influenza. Since I control for district-specific time trends and district-specific monthly effects, all coefficients are identified based on deviations from pre-existing district-specific trends and seasonality patterns. Due to the inclusion of mother fixed effects, all effects are identified based on variation in influenza and center exposure across siblings.

In Equation (2), I make implicit assumptions about the infants. Like Bütikofer et al. (2017), I assume that infants reside in their medical district of birth during their first 12 months of life. Since my data does not include the exact date of birth, I also assume that children were born on the last day of their birth month, being ”treated” by the influenza

prevalence of their birth month and the eight months prior to this month.

For two reasons, the standard errors are clustered at the medical district level. First, individuals within a group may have similar backgrounds and may be exposed to the same environmental factors (Angrist and Pischke 2009). In such cases, there might be substantial correlation between the outcomes within a group, which must be accounted for by adjusting the standard errors. This is particularly important when using regressors with little within-group variation, as the correct standard errors may then be much larger than the conventional ones. I therefore cluster the standard errors at the medical district level, both in order to account for shared background characteristics, and because influenza rates and access to mother and child health care centers vary only at the district level. Since my sample consists of more than 350 medical districts, the consistency of the clustered estimator should not be an issue.

Clustering at the medical district level also solves another problem, namely that of serial correlation in the error terms. With a DiD strategy, shocks that are common to all individuals in given districts and years, such as weather shocks or local reforms, may confound our estimates since it will be hard to distinguish the causal effect of treatment from the effect of the local shocks (Angrist and Pischke 2009). The problem can be mitigated by including multiple periods and/or multiple groups, increasing the chance of such shocks being zero on average. Still, we may expect shocks to be serially correlated across time. The problem can be solved by clustering at a higher level, as long as the higher level includes at least 42 clusters. In my case, since I have a large number of medical districts, clustering at the medical district level mitigates the problem of potential district-year-level shocks.

5 Results

Below, I present the results of my analysis. I start by describing my findings using the entire sample. Then I present results for various subgroups in order to investigate whether there may be heterogeneity by gender, socioeconomic status or geographical area. Lastly, I verify that my results are robust to changing a variety of assumptions.

5.1 Main Results

For education, I look at effects on years of schooling both on the intensive and extensive margin.[25] For earnings, in addition to the log average discounted earnings between 1967 and 2010, I study effects on log earnings in two important periods of life: Between the age of 31 and 40, and between the age of 41 and 50.

Table 2 presents the estimated effect of influenza and access to mother and child health care centers on education and earnings using Equation (2). For earnings, I use logs in order for the linearity assumption to be more closely satisfied, and to make the estimates less sensitive to outliers (Wooldridge 2012). Each column represents a separate regression, with (1) completed years of schooling, (2) whether individuals took more than the mandatory seven years of schooling, (3) log average discounted earnings between 1967 and 2010, (4) log average discounted earnings between the age of 31 and 40, and (5) log average discounted earnings between 41 and 50 as the dependent variables. The dependent variable is regressed on the average number of reported influenza cases per 1,000 inhabitants during the pregnancy as a whole, a dummy indicating whether the infant had access to a mother and child health care center before the age of 1, and an interaction between the two. All specifications include gender, birth order, and dummies indicating

[25] For children in my sample, the mandatory level of schooling was seven years (Store norske leksikon 2015).

Table 2: Long-Term Effects, Pregnancy as a Whole

	Education (margin)		Log Earnings		
	(1)	(2)	(3)	(4)	(5)
	Intensive	Extensive	1967-2010	Age 31-40	Age 41-50
Centers	0.376*	0.034	-0.025	-0.116	-0.082
	(0.161)	(0.028)	(0.071)	(0.081)	(0.098)
Flu	0.062	-0.003	0.011	-0.017	-0.017
	(0.053)	(0.008)	(0.017)	(0.025)	(0.019)
Centers*Flu	-0.019	0.007	-0.003	0.006	0.008
	(0.050)	(0.008)	(0.017)	(0.024)	(0.019)
Observations	38,866	38,866	47,047	44,927	45,549
Adjusted R^2	0.285	0.296	0.380	0.485	0.345

Notes: 'Intensive margin' refers to completed years of schooling. 'Extensive margin' refers to whether individuals took more than the mandatory level of schooling. Each column represents a separate regression where the outcome of interest is regressed on the average influenza prevalence per 1,000 inhabitants during the pregnancy, a dummy indicating whether the infant had access to a mother and child health care center, and an interaction between the two. All specifications include gender, birth order, and dummies indicating whether the mother was unmarried and whether the parents' education levels and age at the time of the child's birth were below average, as well as the ratios of doctors to inhabitants at the time of the child's birth and students to teachers in the year the child was enrolled at school. Additionally, I control for cohort, district and mother fixed effects, as well as a district-specific linear time trend. Effects are therefore identified as deviations from the trend. In order to estimate the effect of influenza rates deviating from the average seasonality in the district, I include month-district fixed effects. The sample includes all non-missing observations for children born between 1935 and 1945 in municipalities that got a mother and child health care center before 1956. Robust standard errors clustered at the level of the medical district of birth are shown in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

whether the mother was unmarried and whether the parents' education levels and age at the time of the child's birth were below average, as well as the ratios of doctors to inhabitants at the time of the child's birth and students to teachers in the year the child was enrolled at school. Additionally, I control for cohort, district and mother fixed effects, as well as a district-specific linear time trend. Effects are therefore identified as deviations from the trend. In order to estimate the effect of influenza rates deviating from the average seasonality in the district, I include month-district fixed effects. The sample includes all non-missing observations for children born between 1935 and 1945 in municipalities that got a mother and child health care center before 1956.

Column (1) of Table 2 shows that access to mother and child health care centers is estimated to increase the years of schooling with somewhat more than one third of a school year.[26] While this is a substantial effect, it does not seem to vary with the average influenza prevalence during pregnancy. Across all specifications, both the effect of influenza and the interaction term between influenza and access to a mother and child health care center are imprecisely measured. In Columns (2)-(5), the coefficient on access to a mother and child health care center is no longer significant, pointing to the effect of well-child visits on the completed years of schooling neither running through the educational extensive margin, nor being mirrored by an increase in earnings. In Section 6, I describe weaknesses with the earnings measures that could potentially explain the counterintuitive finding that earnings is not significantly affected by access to well-child visits, in spite of the large effect on education. However, the effect of well-child visits conditional on no influenza is not the main topic of this thesis.

Based on the literature, I expect long-term educational and labor market outcomes to be affected mainly by second trimester exposure. One possible explanation for the lack of

[26]There are several reasons why the coefficient on access to mother and child health care centers differs from the estimates in Bütikofer et al. (2017). My sample includes data on the centers' opening months. I only use the first center in each medical district, and only individuals born between 1935 and 1945. Lastly, my estimate is conditional on a zero influenza infection rate (Rossin-Slater and Wüst 2016).

significant effects of influenza in the pregnancy as a whole could therefore be that important second trimester effects are hidden by the lack of effects of first and third trimester exposure. In Table 3, I therefore look at effects by trimester of exposure on educational and earnings measures. All specifications are similar to their counterparts in Table 2, except that instead of specifying influenza as the average influenza prevalence per 1,000 inhabitants during the pregnancy months, I include one variable per trimester, each defined as the average influenza prevalence per 1,000 inhabitants during the trimester of interest. Because I continue to include month-district fixed effects, effects are still identified based on deviations from the typical district seasonality.

Since the only difference between Tables 2 and 3 is how the influenza shock is specified, it is reassuring to see that the estimated effect of health-care centers is practically unchanged. Column (1) of Table 3 indicates that access to mother and child health care center increases the number of completed years of schooling, while still not presenting any evidence that the effect depends on whether or not the infant was exposed to influenza as a fetus. None of the outcomes in Columns (2)-(5) seem to be affected by either of the shocks.

Although all trimesters are included to avoid omitted variable bias, I am mainly interested in the effect of exposure in the second trimester. This coefficient is imprecisely estimated in all columns. Surprisingly, the estimated effect of influenza during the first trimester is positive and statistically significant at the five percent level, possibly reflecting a survival bias due to infants surviving first trimester exposure being more able than other infants. Since there is a large degree of negative correlation between exposure in the different trimesters due to the strong seasonality pattern of influenza, a survival bias in any trimester could also confound my coefficients on the second trimester effect. The positive coefficient on exposure in the first trimester may therefore indicate that first trimester selection could be a relevant worry.

Table 3: Long-Term Effects by Trimester

	Education (margin)		Log Earnings		
	(1)	(2)	(3)	(4)	(5)
	Intensive	Extensive	1967-2010	Age 31-40	Age 41-50
Centers	0.373*	0.034	-0.030	-0.117	-0.029
	(0.162)	(0.028)	(0.071)	(0.081)	(0.018)
Flu trimester 1	0.062*	0.001	0.016	0.004	-0.003
	(0.027)	(0.004)	(0.009)	(0.015)	(0.003)
Flu trimester 2	0.010	0.001	0.013	-0.014	0.001
	(0.025)	(0.004)	(0.011)	(0.012)	(0.002)
Flu trimester 3	-0.004	-0.006	-0.015	-0.005	-0.003
	(0.033)	(0.006)	(0.010)	(0.013)	(0.002)
Centers*Flu trimester 1	-0.038	-0.001	0.004	0.001	0.016
	(0.034)	(0.004)	(0.012)	(0.018)	(0.011)
Centers*Flu trimester 2	0.010	0.003	-0.020	0.004	-0.002
	(0.034)	(0.005)	(0.013)	(0.013)	(0.013)
Centers*Flu trimester 3	0.004	0.004	0.015	0.001	-0.003
	(0.033)	(0.007)	(0.011)	(0.016)	(0.011)
Observations	38,866	38,866	47,047	44,927	45,549
Adjusted R^2	0.286	0.296	0.381	0.485	0.346

Notes: 'Intensive margin' refers to completed years of schooling. 'Extensive margin' refers to whether individuals took more than the mandatory level of schooling. Each column represents a different regression of the variable indicated by the column name on the influenza prevalence per 1,000 inhabitants during each trimester, a dummy indicating access to a mother and child health care center, and an interaction between each trimester and the center dummy. All specifications include gender, birth order, and dummies indicating whether the mother was unmarried and whether the parents' education levels and age at the time of the child's birth were below average, as well as the ratios of doctors to inhabitants at the time of the child's birth and students to teachers in the year the child was enrolled at school. Additionally, I control for cohort, district and mother fixed effects, as well as a district-specific linear time trend. Effects are therefore identified as deviations from the trend. In order to estimate the effect of influenza rates deviating from the average seasonality in the district, I include month-district fixed effects. Robust standard errors clustered at the level of the medical district of birth are shown in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

It is important to understand the extent to which such survival bias may confound my estimates. I therefore look closer at selection effects over the life cycle. I start by analyzing how influenza affects the rates of stillbirths and infant mortality. Since I have only yearly observations of these outcomes, I generate yearly influenza measures. From the individual-level analysis, for every birth month I have measures of the average influenza prevalence per 1,000 inhabitants during the pregnancy period. I therefore define the yearly measures as the average of the monthly measures. This approach has the benefit of taking all relevant months into account, also those prior to the year of interest, and weighting the importance of each monthly infection rate by how many of the infants born in the given year would have been in utero during that period. For instance, the infection rate in a month like January, which would have affected a large share of the infants born that year, is given a larger weight in the calculation of the yearly measure than the infection rate in December, which only affected the infants born that month.

In both the stillbirths and the infant mortality analyses, I control for year and district fixed effects. In the stillbirth analysis, I additionally include a district-specific trend. I run regressions by gender because male infants may be adversely affected by negative in utero shocks (see for instance Kraemer 2000). The estimated effects are presented in the four first columns of Table A2. Influenza does not seem to increase the relative number of stillbirths or infant deaths.[27]

I then turn to long-term selection effects, following Schwandt (2017) on studying whether third trimester exposure increases the fraction of missing outcomes in adulthood.[28] Columns (5)-(6) of Table A2 presents the estimated long-run selection effects. Since my sample consists of individuals who were still alive in 1967, the number of individuals who do not show up in the education or earnings data will necessarily be small. I therefore do not

[27] Replacing the bottom-coded stillbirth observations with a randomly assigned number below three dramatically increases the sample, while still yielding similar patterns.

[28] All trimesters are included in the analysis to avoid omitted variable bias due to the high correlation between the influenza prevalence in the different trimesters.

expect to find large effects. Nevertheless, it is reassuring to see that there is no evidence of third trimester exposure to maternal influenza substantially increasing the number of individuals who are not observable in adulthood.[29] In sum, although the statistically significant positive effect of first trimester exposure to influenza could indicate a survival bias, selection effects do not overall seem to drive my results.

5.2 Heterogeneous Effects

In the last section, I considered short-run impacts of influenza separately by gender. There are several reasons to look at possible heterogeneous effects also for the long-run outcomes. As pointed out by Bütikofer et al. (2017), girls born in the period I study had different labor market chances than boys. In addition, both of the shocks I study have been shown to have different impacts on boys and girls. Bütikofer et al. (2017) find that the effect of mother and child health care centers on earnings and years of education was larger for boys.[30] Schwandt (2017) finds that influenza affected wages and labor force participation more strongly for male fetuses, while the effect on welfare rates was larger for females. Based on my main specification, I run a new regression where for all independent variables of interest, I include an additional interaction term between the variable and a dummy indicating whether the infant was female.

Table 4 presents the results of the analysis by gender. 'Education' here refers to the years of completed schooling, whereas 'Log earnings' refers to the log average discounted earnings between 1967 and 2010. For years of schooling, there is no evidence that the interaction between influenza and mother and child health care centers depends on the gender of the child. For earnings, the access to mother and child health care center depends more strongly on exposure to maternal influenza in the second trimester if the child is female.

[29] There was not enough variation in the data to estimate this regression separately for each gender.
[30] They cannot reject the hypothesis that the effect sizes are the same.

Table 4: Second Trimester Effects by Gender

	(1) Education	(2) Log Earnings
Centers*Flu trimester 2	0.016	0.008
	(0.046)	(0.013)
Centers*Flu trimester 2*Female	-0.021	-0.065**
	(0.058)	(0.021)

Notes: 'Education' here refers to the years of completed schooling, whereas 'Log earnings' refers to the log average discounted earnings between 1967 and 2010. Each column represents a different regression of the variable indicated by the column name on the influenza prevalence per 1,000 inhabitants during each trimester, a dummy indicating access to a mother and child health care center, and an interaction between each trimester and the center dummy. In addition, for all independent variables of interest an interaction term between the variable and a dummy indicating whether the child is female is included. All specifications include gender, birth order, and dummies indicating whether the mother was unmarried and whether the parents' education levels and age at the time of the child's birth were below average, as well as the ratios of doctors to inhabitants at the time of the child's birth and students to teachers in the year the child was enrolled at school. Additionally, I control for cohort, district and mother fixed effects, as well as a district-specific linear time trend. Effects are therefore identified as deviations from the trend. In order to estimate the effect of influenza rates deviating from the average seasonality in the district, I include month-district fixed effects. Robust standard errors clustered at the level of the medical district of birth are shown in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Given the low level of female labor force participation at the time, this effect may however be identified based on a selected subgroup. I therefore mainly view Table 4 as reassuring evidence that the lack of significant interaction effect does not hide substantial effects for males.

The effects of influenza and well-child visits could also differ with socioeconomic status (Bütikofer et al. 2017; Schwandt 2017). While I cannot observe the income of the parents in my sample, I can observe which mothers were unmarried, which should be highly associated with low-income status. However, the external validity of this analysis may be rather low, as unmarried mothers is likely to be a highly selected group. As an alternative

proxy for low-income status, I therefore look at whether fathers had less than average levels of education.

In Table 5, I report the estimated interaction term between second trimester influenza exposure and mother and child health care center for the subsample of infants whose mothers were not married, and the subsample of infants whose fathers had lower than average levels of education. Again, I include the results using the entire sample for comparison purposes.

The effect presented in Column (3) of Table 5 indicates that for children of unmarried mothers, the effect of mother and child health care centers on years of schooling is significantly lower if the child was exposed to maternal influenza in the second trimester. However, the magnitude of the estimated coefficient seems implausible. Specifically, increasing the influenza prevalence with 1 per 1,000 inhabitants seems to increase the impact of centers with more than one third of a school year. Given that Figure 1 revealed that the influenza prevalence per 1,000 inhabitants gets as high as 10 in my sample, the estimated interaction effect would suggest that mother and child health care centers increase years of schooling with multiple years for those most exposed to influenza in the second trimester. A possible explanation is that for the sample of unmarried mothers, mother-fixed effects may greatly overstate the extent to which children are treated the same. One sibling being ”treated” for instance by adoptive parents with high socioeconomic status could bias my results, particularly due to the small sample size. In Section 6, I return to the issue of reinforcing parental behavior in sibling fixed effects models. In sum, although Columns (3)-(4) provide interesting information, it does not seem like the estimated interaction term between influenza and well-child visits for children born to unmarried mothers is representative for poor children as a whole.

Columns (5)-(6) present the estimated effects for the subsample of infants born to fathers with lower than average levels of education, which should be more representative for the

population of poor children. Here, I find no evidence that the effect of well-child visits on earnings or education depends on influenza exposure. This seems to suggest that the effect of influenza on mother and child health care centers does not depend on socioeconomic status in general.

Table 5: Second Trimester Effects by Parental Socioeconomic Status

	Full sample		Mother unmarried		Father low education	
	(1)	(2)	(3)	(4)	(5)	(6)
	Education	Log Earnings	Education	Log Earnings	Education	Log Earnings
Centers*Flu trimester 2	0.010	-0.020	-0.357*	-0.050	-0.062	-0.047
	(0.034)	(0.013)	(0.162)	(0.138)	(0.056)	(0.024)
Observations	38,866	47,047	3,722	4,800	17,581	21,370
Adjusted R^2	0.286	0.381	0.922	0.974	0.486	0.518

Notes: The first title row indicates which subsample is being used. 'Education' here refers to the years of completed schooling, whereas 'Log earnings' refers to the log average discounted earnings between 1967 and 2010. Each column represents a different regression of the variable indicated by the column name on the influenza prevalence per 1,000 inhabitants during each trimester, a dummy indicating access to a mother and child health care center, and an interaction between each trimester and the center dummy. Only the coefficient on the interaction effect for the second trimester is reported. All specifications include gender, birth order, and dummies indicating whether the mother was unmarried and whether the parents' education levels and age at the time of the child's birth were below average, as well as the ratios of doctors to inhabitants at the time of the child's birth and students to teachers in the year the child was enrolled at school. Additionally, I control for cohort, district and mother fixed effects, as well as a district-specific linear time trend. Effects are therefore identified as deviations from the trend. In order to estimate the effect of influenza rates deviating from the average seasonality in the district, I include month-district fixed effects. Robust standard errors clustered at the level of the medical district of birth are shown in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

In the last part of my heterogeneity analysis, I ask whether there is geographical variation in the estimated effects. In particular, at this time the northern parts of the country differed from the rest in terms of ethnic composition of the population, especially due to the Sami and Kven population. There could therefore be both biological and cultural reasons why the effects of the two shocks I study could be different in these parts of the country. Additionally, parts of the northernmost counties are very sparsely populated, which could potentially affect the spread of both influenza viruses and information related to infant nutrition and hygiene. For every part of the country, I therefore run a separate regression where I include interaction terms between a dummy describing the geographical location of the medical district of birth, and the independent variables. In Table 6, for every part of the country I report the estimated interaction effect between the location dummy and the main variable of interest, which is the interaction term between second trimester exposure to maternal flu and mother and child health care centers. It does not seem like the effect of influenza on the returns to mother and child health care centers differs with geographical location.

5.3 Sensitivity Analysis

In this subsection, I present a variety of robustness checks. First, I test how my results are affected by changing the assumptions I make in my main analysis. Then, I verify that my results are robust to excluding selected observations that may differ from the rest of the sample. I also check that the lack of significant effects of mother and child health care centers is not due to specifying treatment on the medical district level instead of the municipality level. Lastly, I show that neither influenza nor access to mother and child health care centers seems to affect education or earnings in my sample, and that my analysis is not likely to suffer from a problem with 'bad controls'.

Table 6: Second Trimester Effects by Geographical Location

	(1) Education	(2) Log Earnings
Nordland, Troms and Finnmark	-0.101	-0.015
	(0.133)	(0.029)
Sør- and Nord-Trøndelag	-0.071	-0.004
	(0.043)	(0.032)
Rogaland, Hordaland, Sogn og Fjordane and Møre og Romsdal	0.052	0.038
	(0.067)	(0.038)
Vestfold, Telemark, Aust-Agder and Vest-Agder	-0.174	-0.035
	(0.114)	(0.066)
Østfold, Akershus, Hedmark, Oppland and Buskerud	0.036	-0.007
	(0.062)	(0.023)

Notes: 'Education' here refers to the years of completed schooling, whereas 'Log earnings' refers to the log average discounted earnings between 1967 and 2010. Each cell represents a different regression of the variable indicated by the column name on the influenza prevalence per 1,000 inhabitants during each trimester, a dummy indicating access to a mother and child health care center, and an interaction between each trimester and the center dummy. In addition, for all variables of interest, an interaction term between the variable and a dummy indicating the geographical location of the medical district is included. Only the interaction term between the geographical indicator, second trimester influenza prevalence and the center dummy is reported. All specifications include gender, birth order, and dummies indicating whether the mother was unmarried and whether the parents' education levels and age at the time of the child's birth were below average, as well as the ratios of doctors to inhabitants at the time of the child's birth and students to teachers in the year the child was enrolled at school. Additionally, I control for cohort, district and mother fixed effects, as well as a district-specific linear time trend. Effects are therefore identified as deviations from the trend. In order to estimate the effect of influenza rates deviating from the average seasonality in the district, I include month-district fixed effects. Robust standard errors clustered at the level of the medical district of birth are shown in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 7 shows how my results are affected by changing various assumptions. I use the same specifications as those presented in Table 3, including the influenza prevalence in all trimesters. To avoid clutter, I only report the coefficient of main interest, that is, how the effect of mother and child health care centers is affected by the influenza exposure in the second trimester. The row name describes which assumption is changed.

In the analysis presented in the first row of Table 7, I change the assumption of the exact day of birth. By assuming that the infant was born on the first day of the month instead of the last, I implicitly assume that the infant was "treated" by slightly different influenza rates than in my main specifications. In the second row, I assume that mother and child health care centers for which I do not know the opening month, opened in July instead of January. In the third and fourth rows, since all effects are identified based on deviations from a time trend, I verify that results are not driven by how this trend is specified. Instead of using the linear time trend as in my main specifications, I specify the time trend as quadratic and cubic, respectively. None of these changes in assumptions seem to drive the lack of significant interaction effects of second trimester exposure to influenza and access to well-child visits on education or earnings.

In the fifth row of Table 7, I look into how my results are affected by controlling for month fixed effects instead of month-district fixed effects. This is the approach Schwandt (2017) uses in his cohort analysis, finding a significant and negative impact of second trimester exposure to maternal influenza on years of schooling. I therefore include this specification to see if the lack of a significant interaction effect between the second trimester exposure and well-child visits may be due to my approach reducing the amount of variation in my data. This does not seem to be the case. None of the estimated interaction effects are significant even when specifying the seasonality pattern on a national instead of a local level.

The sixth row of Table 7 presents the result of an alternative approach to defining which

medical districts are treated by mother and child health care center. As explained in Section 3, in my main analysis I include only the municipalities that received a mother and child health care center before 1956. By not excluding medical districts with non-treated municipalities altogether, this approach could potentially attenuate my results, as I include districts that could potentially have been treated by centers not run by NKS before I observe the first NKS center. In the last row, I therefore look into how my results are affected by excluding all districts that contain at least one municipality that did not get an NKS mother and child health care center before 1956.

For earnings, another possibility could be that only the upper tail of the earnings distribution is affected by well-child visits. In this case, my earnings measures would be poorly suited to capture the effect due to the logs decreasing the amount of variation in the upper tail of the measures. In Table A3, I look into whether the two shocks affect earnings when measured in levels instead of logs. The specification is the same as before, with influenza prevalence specified per trimester. I find no evidence of interaction effects between access to well-child visits and second trimester influenza exposure, which suggests that the lack of effect on earnings is not due to effects in the upper tail of the earnings distribution being muted by the log measure.

I then turn to excluding selected parts of the sample in order to verify that particular observations do not drive my results. In Table 8, each cell represents a different regression, with the column name indicating which dependent variable is being used, and the row name describes which part of the sample has been excluded. The specification is unchanged from the main analysis, still including all three trimesters. I continue to report the estimated interaction effect between influenza in the second trimester and access to mother and child health care center only.

As explained in Section 3, the influenza prevalence in Southern Norway may be overstated for the fall of 1944 and 1945. In the first row of Table 8, I therefore present evidence

Table 7: Basic Robustness Tests

	Education (margin)		Log Earnings		
	(1)	(2)	(3)	(4)	(5)
	Intensive	Extensive	1967-2010	Age 31-40	Age 41-50
July as missing opening month	0.009	0.003	-0.020	0.004	0.001
	(0.035)	(0.005)	(0.013)	(0.013)	(0.003)
Observations	38,866	38,866	47,047	44,927	45,549
Adjusted R^2	0.286	0.296	0.381	0.485	0.346
Born on first day of month	0.044	0.002	-0.022	0.005	0.003
	(0.027)	(0.004)	(0.014)	(0.012)	(0.013)
Observations	38,866	38,866	46,854	44,927	45,360
Adjusted R^2	0.286	0.296	0.382	0.485	0.348
Quadratic trend	0.009	0.002	-0.020	0.003	-0.001
	(0.034)	(0.005)	(0.013)	(0.013)	(0.013)
Observations	38,866	38,866	47,047	44,927	45,549
Adjusted R^2	0.285	0.292	0.381	0.485	0.346
Cubic trend	0.007	0.001	-0.019	0.003	-0.001
	(0.034)	(0.005)	(0.013)	(0.013)	(0.013)
Observations	38,866	38,866	47,047	44,927	45,549
Adjusted R^2	0.285	0.289	0.381	0.484	0.345
National seasonality pattern	0.009	0.005	-0.019	0.001	-0.007
	(0.026)	(0.003)	(0.010)	(0.012)	(0.011)
Observations	38,866	38,866	47,047	44,927	45,549
Adjusted R^2	0.089	0.107	0.288	0.383	0.202
Re-defined center treatment	0.021	0.003	-0.010	0.007	0.004
	(0.042)	(0.006)	(0.014)	(0.018)	(0.013)
Observations	33,075	33,075	40,104	38,301	38,799
Adjusted R^2	0.269	0.290	0.366	0.475	0.330

Notes: 'Intensive margin' refers to completed years of schooling. 'Extensive margin' refers to whether individuals took more than the mandatory level of schooling. Each cell represents a different regression of the variable indicated by the column name on the influenza prevalence per 1,000 inhabitants during each trimester, a dummy indicating access to a mother and child health care center, and an interaction between each trimester and the center dummy. Only the coefficient on the interaction effect for the second trimester is reported. All specifications include gender, birth order, and dummies indicating whether the mother was unmarried and whether the parents' education levels and age at the time of the child's birth were below average, as well as the ratios of doctors to inhabitants at the time of the child's birth and students to teachers in the year the child was enrolled at school. Additionally, I control for cohort, district and mother fixed effects, as well as a district-specific time trend, which is specified as linear unless otherwise indicated. Effects are therefore identified as deviations from the trend. In order to estimate the effect of influenza rates deviating from the average seasonality in the district, I include month-district fixed effects. The sample includes all non-missing observations for children born between 1935 and 1945 in municipalities that got a mother and child health care center before 1956. The row name describes which assumption is changed. Robust standard errors clustered at the level of the medical district of birth are shown in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

that the estimated interaction effect does not substantially change when excluding the infants born after the northern parts of the country were evacuated.[31] In the second row, I present a specification where I exclude all births during World War II (WWII).[32] When excluding infants born after April 1940, I find some evidence of a positive interaction effect between influenza in the second trimester and access to mother and child health care centers, pointing to centers increasing the years of schooling more for those affected by influenza. The relationship could reflect differences between the pre-war and the war period, or noise in my data that is given more weight when reducing the sample with more than 50 percent. However, the effect of well-child visits is not significantly different from zero in this setup, making it likely that the interaction effect is in part picking up the effect of mother and child health care centers.

In rows 3-6 of Table 8, I verify that my results are not driven by outliers. I start by excluding observations at the 1st and 99th percentile of the average influenza prevalence during the pregnancy as a whole.[33] These results are presented in row 3 and 4. In row 5 and 6, for each outcome I remove the 1st and 99th percentile respectively, before re-running my analysis. Every column therefore represents a separate subsample. Note that these cells are empty for the column presenting the effect of the two shocks on the educational extensive margin. Since the dependent variable is a dummy variable, excluding outliers makes no sense. None of the reported interaction terms between influenza in the second trimester and mother and child health care centers are significant, indicating that the lack

[31] Both for this analysis and the analysis where I exclude the entire war period, I exclude the entire year of 1945 since the occupation likely had long-lasting consequences.

[32] Norway was occupied by Nazi Germany from April 1940 to May 1945, but the rollout of mother and child health care centers continued during the war (Bütikofer et al. 2017). Whether food rationing improved or worsened health during this period is still debated. Rationing of goods like sugar, tobacco and alcohol may have contributed to improved health, whereas food scarcity and unbalanced nutrition seem to have left the population less resistant to infectious diseases such as tuberculosis and scarlet fever (Njølstad 2015).

[33] Using the 1st and 99th percentile as measures of outliers has important limitations, as the percentiles are directly affected by outliers themselves. However, since this test is not central to my analysis, but merely serves to check that my results are not driven by the most extreme observations, I use the percentile measures to keep the interpretation simple.

of significant interaction effects in my main analysis is not driven by outliers.

Since the coefficients on influenza and access to mother and child health care center are mostly not significantly different from zero in my sample, it is interesting to look at whether the shocks had a significant impact on their own. In Table A4, I present the results of regressions that only include influenza in the different trimesters as regressors. The control variables are the same as in my main specifications. The coefficient on influenza exposure in the second trimester is not significant when looking only at the influenza shock either.

Table A5 presents similar regressions using only the well-child visits shock as the regressor. It does not seem like access to mother and child health care center has any detectable effects on education on the extensive margin or on earnings in my sample. Since this result is inconsistent with the estimated impact of well-child visits on earnings in Bütikofer et al. (2017), I replicate the analysis from Bütikofer et al. (2017) on the cohorts in my sample to verify that the difference in effect size does not stem from specifying the center treatment on a medical district level instead of on the municipality level. I find no evidence that the earnings of individuals born between 1935 and 1945 were affected by access to centers, even when treatment is defined on a municipality level.[34] Effects of mother and child health care centers therefore seem to be cohort-specific, with less impact for the cohorts I study. In sum, although I find little evidence of dynamic complementarities, it may very well be that the lack of significant interaction effects stems from the lack of impact of the shocks on their own.

Separately estimating the impacts of the two shocks also provides useful information about my main specification. One could ask whether influenza rates could potentially be affected by a mother and child health care center, for instance because hygiene would improve or because the disease would spread faster due to mothers meeting at the center.

[34]This table is available on request.

Table 8: Removing Selected Parts of the Sample

	Education (margin)		Log Earnings		
	(1)	(2)	(3)	(4)	(5)
	Intensive	Extensive	1967-2010	Age 31-40	Age 41-50
Removing evacuation period	0.030	0.009	-0.012	0.001	0.006
	(0.064)	(0.008)	(0.018)	(0.023)	(0.017)
Observations	24,161	24,161	30,077	28,346	29,137
Adjusted R^2	0.451	0.458	0.483	0.602	0.465
Removing war period	0.521*	0.057	-0.024	-0.051	-0.066
	(0.232)	(0.052)	(0.069)	(0.102)	(0.065)
Observations	9,897	9,897	13,180	11,931	12,756
Adjusted R^2	0.845	0.859	0.811	0.890	0.835
Removing 1st percentile flu	0.015	0.002	-0.014	0.015	-0.000
	(0.038)	(0.005)	(0.012)	(0.015)	(0.013)
Observations	35,620	35,620	43,166	41,222	41,791
Adjusted R^2	0.322	0.308	0.402	0.503	0.373
Removing 99th percentile flu	0.010	0.003	-0.020	0.004	-0.002
	(0.034)	(0.005)	(0.013)	(0.013)	(0.013)
Observations	38,860	38,860	47,040	44,920	45,543
Adjusted R^2	0.286	0.296	0.381	0.485	0.346
Removing 1st percentile outcome	-0.005	-	-0.020	0.004	-0.002
	(0.044)	-	(0.013)	(0.013)	(0.013)
Observations	33,516	-	47,047	44,926	45,549
Adjusted R^2	0.345	-	0.381	0.485	0.346
Removing 99th percentile outcome	0.009	-	-0.020	0.004	-0.002
	(0.034)	-	(0.013)	(0.013)	(0.013)
Observations	37,723	-	47,046	44,926	45,549
Adjusted R^2	0.291	-	0.381	0.485	0.346

Notes: 'Intensive margin' refers to completed years of schooling. 'Extensive margin' refers to whether individuals took more than the mandatory level of schooling. Each cell represents a different regression of the variable indicated by the column name on the influenza prevalence per 1,000 inhabitants during each trimester, a dummy indicating access to a mother and child health care center, and an interaction between each trimester and the center dummy. Only the coefficient on the interaction effect for the second trimester is reported. The row name describes which observations are excluded. Except from these excluded observations, the sample includes all non-missing observations for children born between 1935 and 1945 in municipalities that got a mother and child health care center before 1956. All specifications include gender, birth order, and dummies indicating whether the mother was unmarried and whether the parents' education levels and age at the time of the child's birth were below average, as well as the ratios of doctors to inhabitants at the time of the child's birth and students to teachers in the year the child was enrolled at school. Additionally, I control for cohort, district and mother fixed effects, as well as a district-specific linear time trend. Effects are therefore identified as deviations from the trend. In order to estimate the effect of influenza rates deviating from the average seasonality in the district, I include month-district fixed effects. Robust standard errors clustered at the level of the medical district of birth are shown in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Such variables that are themselves outcomes are what Angrist and Pischke (2009) call 'bad controls' and should not be included as independent variables. It is therefore reassuring to see that the estimated impact of access to mother and child health care center is practically the same regardless of whether I look at the effect of the centers only, as in Table A5, or the effect of centers holding influenza rates constant, as in Table 3. This points to center openings not having any major impact on the influenza prevalence in a medical district.

Similarly, one could imagine that a center opening would increase the number of doctors in a medical district, thereby increasing the ratio of doctors to inhabitants. In a separate regression, I therefore drop the ratio of doctors to inhabitants from the control variables to see whether the estimated impact of influenza and well-child visits changes. The results are presented in Table A6 and are practically indistinguishable from those presented in Table 3, which are based on specifications including the ratio of doctors to inhabitants. In sum, the results presented in Tables A5 and A6 suggest that influenza and the ratio of doctors to inhabitants should not be 'bad controls'.

6 Discussion

In this section, I discuss possible limitations of my study that are not already covered in the previous sections. Based on these limitations, I provide suggestions for further research.

This thesis finds no evidence of dynamic complementarities. However, important parameters of the human capital production function such as parental investments are unobservable to me (Rossin-Slater and Wüst 2016). Therefore, it is relevant to look at how parental responses could affect my findings. Since all my long-term specifications include sibling fixed effects, compensating parental investments would result in a downward bias,

so that the estimated effects would be too low (Almond and Currie 2011). Reinforcing parental investments, however, could lead me to overstate the true effects of influenza and well-child visits, since the difference between siblings exposed to differential influenza and well-child visits treatment would appear larger than the isolated causal effect of the shocks. Reinforcing parental behavior would therefore be a larger issue.

The literature on parental investments is still rather small and cannot give a conclusive answer as to whether parental investments would most likely have been compensating or reinforcing in the setting I study. However, so far there is little evidence of reinforcing parental behavior in developed countries (Almond and Currie 2011). For my sample, compensating parental behavior seems more likely than reinforcing parental behavior, so that rather than overstating the true effects of the two shocks I consider, my analysis could have understated their impact.

There are several other possible explanations for the lack of significant interaction effects in my analysis. Maybe the most likely explanation is the lack of significant impacts of the two shocks on their own. Since Bütikofer et al. (2017) find an effect of well-child visits using a sample that includes cohorts born after 1945, future research should include influenza data for these cohorts in order to get further in analyzing whether the returns to mother and child health care centers are affected by exposure to maternal influenza.[35] In addition, further research should look at labor earnings and governmental transfers separately by including more detailed earnings measures. Since my earnings measures include both types of earnings, the estimated impact of influenza on earnings could be attenuated by influenza exposure leading to increased governmental transfers (see Schwandt 2017). Lastly, further research should attempt to control for first-trimester abortions as well as premature births in order to improve the precision of the analysis.

[35] Beyond the scope of this thesis, but very interesting, is the question of why well-child visits seem to affect some cohorts more than others.

7 Conclusion

A large literature documents that early-life shocks is important for the production of human capital. However, we do not know much about how different shocks interact. In theoretical work, the idea of dynamic complementarities has been put forward, implying that the effect of a positive shock is larger for individuals with a higher baseline level of skills. In this setting, universal interventions may in reality serve to strengthen initial differences. The policy relevance of understanding whether dynamic complementarities exist can therefore hardly be underestimated, but no clear answer has yet emerged from the literature.

I utilize two arguably exogenous sources of variation in infant health in Norway between 1935 and 1945 to shed light on how the shocks interacted. The first shock is fetal exposure to seasonal influenza, which has been shown to negatively affect long-term health and labor market outcomes. The second shock is access to well-child visits, which is known to improve similar outcomes. Specifically, I ask whether fetal exposure to influenza reduces the long-term returns to mother and child health care centers.

Using detailed individual-level registry data, I find little evidence of influenza affecting the educational or labor market returns to well-child visits. However, the lack of significant interaction effects could stem from the two shocks not significantly affecting the outcomes of interest in the period I study. Future research should therefore collect more detailed data, and data for a longer period of time, in order to understand whether the returns to mother and child health care centers are lower for infants who were exposed to maternal influenza.

References

Adhvaryu, A., T. Molina, A. Nyshadham, and J. Tamayo (2016): "Helping Children Catch Up: Early Life Shocks and the Progresa Experiment". Working paper.

Aizer, A. and F. Cunha (2012): "The Production of Child Human Capital: Endowments, Investments, and Fertility". Working Paper 18429, National Bureau of Economic Research.

Almond, D. (2006): "Is the 1918 Influenza Pandemic Over? Long-Term Effects of In Utero Influenza Exposure in the Post-1940 U.S. Population". *Journal of Political Economy* 114(4), 672–712.

Almond, D. and J. Currie (2011): "Human Capital Development before Age Five". In: *Handbook of Labor Economics*. Ed. by D. Card and O. Ashenfelter. 1st ed. Vol. 4.b. Elsevier. Chap. 15, 1315–1486.

Almond, D., J. Currie, and V. Duque (2017): "Childhood Circumstances and Adult Outcomes: Act II". Working Paper 23017, National Bureau of Economic Research.

Almond, D. and B. Mazumder (2011): "Health Capital and the Prenatal Environment: The Effect of Ramadan Observance during Pregnancy". *American Economic Journal: Applied Economics* 3(4), 56–85.

Almond, D. and B. Mazumder (2013): "Fetal Origins and Parental Responses". *Annual Review of Economics* 5, 37–56.

Angrist, J. D. and J.-S. Pischke (2009): *Mostly Harmless Econometrics. An Empiricist's Companion*. Princeton, NJ: Princeton University Press.

Bakken, M. R. and S. J. Husøy (2016): "The Long Term Impact of the 1918 Influenza Pandemic in Norway". Norwegian School of Economics.

Bhalotra, S. R., M. Karlsson, and T. Nilsson (2016): "Infant Health and Longevity: Evidence from a Historical Trial in Sweden". Working paper.

Bharadwaj, P., J. V. Johnsen, and K. V. Løken (2014): "Smoking Bans, Maternal Smoking and Birth Outcomes". *Journal of Public Economics* 115, 72–93.

Bleakley, H. (2007): "Disease and Development: Evidence from Hookworm Eradication in the American South". *Quarterly Journal of Economics* 122(1), 73–117.

Bloom-Feshbach, K., L. Simonsen, C. Viboud, K. Mølbak, M. A. Miller, M. Gottfredsson, and V. Andreasen (2011): "Natality Decline and Miscarriages Associated With the 1918 Influenza Pandemic: The Scandinavian and United States Experiences". *Journal of Infectious Diseases* 204(8), 1157–1164.

Brown, R. and D. Thomas (2016): "On the Long Term Effects of the 1918 U.S. Influenza Pandemic". Working paper.

Bütikofer, A., K. V. Løken, and K. G. Salvanes (2017): "Long-Term Consequences of Access to Well-child Visits". Working paper.

Carneiro, P., K. V. Løken, and K. G. Salvanes (2015): "A Flying Start? Maternity Leave Benefits and Long-Run Outcomes of Children". *Journal of Political Economy* 123(2), 365–412.

NCIRD – Centers for Disease Control and Prevention, National Center for Immunization and Respiratory Diseases (2016): *Key Facts About Influenza (Flu)*. `www.cdc.gov/flu/keyfacts.htm` (accessed: 2017–02–06).

NCIRD – Centers for Disease Control and Prevention, National Center for Immunization and Respiratory Diseases (2017): *Pregnant Women & Influenza (Flu)*. `www.cdc.gov/flu/protect/vaccine/pregnant.htm` (accessed: 2017–02–06).

Chen, A., E. Oster, and H. Williams (2016): "Why Is Infant Mortality Higher in the United States Than in Europe?" *American Economic Journal: Economic Policy* 8(2), 89–124.

Cunha, F. and J. J. Heckman (2007): "The Technology of Skill Formation". *American Economic Review* 97(2), 31–47.

Currie, J. and H. Schwandt (2013): "Within-mother analysis of seasonal patterns in health at birth". *Proceedings of the National Academy of Sciences* 110(28), 12265–12270.

Currie, J. and R. Walker (2011): "Traffic Congestion and Infant Health: Evidence from E-ZPass". *American Economic Journal: Applied Economics* 3, 65–90.

Deaton, A. (2010): "Instruments, Randomization, and Learning about Development". *Journal of Economic Literature* 48, 424–455.

Duque, V., M. Rosales-Rueda, and F. Sanchez (2016): "Integrating Early-Life Shocks and Human Capital Investments on Children's Education". Working paper 2016/06, CAF (Development Bank of Latin America).

Galiani, S., P. Gertler, and E. Schargrodsky (2005): "Water for Life: The Impact of the Privatization of Water Services on Child Mortality". *Journal of Political Economy* 113(1), 83–120.

Garces, E., D. Thomas, and J. Currie (2002): "Longer-Term Effects of Head Start". *American Economic Review* 92(4), 999–1012.

Gilraine, M. (2016): "School Accountability and the Dynamics of Human Capital Formation". Working paper.

Gunnsteinsson, S., A. Adhvaryu, P. Christian, A. Labrique, J. Sugimoto, A. A. Shamim, and K. P. W. Jr. (2016): "Evidence From the Interaction of a Natural Experiment and a Randomized Control Trial". Working paper.

Havnes, T. and M. Mogstad (2011): "No Child Left Behind: Subsidized Child Care and Children's Long-Run Outcomes". *American Economic Journal: Economic Policy* 3, 97–129.

Hjort, J., M. Sølvsten, and M. Wüst (forthcoming): "Universal Investment in Infants and Long-Run Health: Evidence from Denmark's 1937 Home Visiting Program". *American Economic Journal: Applied Economics*.

Johnson, R. C. and C. K. Jackson (2017): "Reducing Inequality Through Dynamic Complementarity: Evidence from Head Start and Public School Spending". Working Paper.

Juvkam, D. (1999): *Historisk oversikt over endringer i kommune- og fylkesinndelingen*. Oslo/Kongsvinger: Statistics Norway. `www.ssb.no/a/publikasjoner/pdf/rapp_9913/rapp_9913.pdf` (accessed: 2017–03–23).

Kelly, E. (2011): "The Scourge of Asian Flu: In utero Exposure to Pandemic Influenza and the Development of a Cohort of British Children". *Journal of Human Resources* 46(4), 669–694.

Kraemer, S. (2000): "The fragile male". *BMJ* 321(7276), 1609–1612.

Lin, M.-J. and E. M. Liu (2014): "Does in utero exposure to Illness matter? The 1918 influenza epidemic in Taiwan as a natural experiment". *Journal of Health Economics* 37, 152–163.

Malamud, O., C. Pop-Eleches, and M. Urquiola (2016): "Interactions Between Family and School Environments: Evidence on Dynamic Complementarities?" Working Paper 22112, National Bureau of Economic Research.

Moehling, C. M. and M. A. Thomasson (2014): "Saving Babies: The Impact of Public Education Programs on Infant Mortality". *Demography* 51 (2), 367–386.

Neelsen, S. and T. Stratmann (2012): "Long-run effects of fetal influenza exposure: Evidence from Switzerland". *Social Science & Medicine* 74 (1), 58–66.

Nelson, R. E. (2010): "Testing the Fetal Origins Hypothesis in a developing country: evidence from the 1918 Influenza Pandemic". *Health Economics* 19(10), 1181–1192.

Njølstad, O. (2015): *Kosthold og ernæring under okkupasjonen*. University of Oslo. www.norgeshistorie.no/andre-verdenskrig/hus-og-hjem/1704-kosthold-og-ernering-under-okkupasjonen.html (accessed: 2017–06–10).

Petterson, A. (2016): *Tvangsevakuering og overvintring i Finnmark og Nord-Troms 1944–1945*. University of Oslo. www.norgeshistorie.no/andre-verdenskrig/artikler/1736-tvangsevakuering-og-overvinting-i-finnmark-og-nord-troms-1944-45.html (accessed: 2017–04–25).

Richter, A. and P. O. Robling (2016): "Multigenerational effects of the 1918-19 influenza pandemic on educational attainment: Evidence from Sweden". Working Paper.

Rossin-Slater, M. and M. Wüst (2016): "What is the Added Value of Preschool? Long-term Impacts and Interactions with a Health Intervention". Working Paper 22700, National Bureau of Economic Research.

Schiøtz, A. (2003): *Folkets helse - landets styrke 1850-2003*. Oslo: Universitetsforlaget.

Schwandt, H. (2017): "The Lasting Legacy of Seasonal Influenza: In-utero Exposure and Labor Market Outcomes". Working Paper.

Statistics Norway (1932): *Norges civile, geistlige, rettslige og militære inndeling 1. juli 1932*. Oslo. www.ssb.no/a/histstat/nos/nos_viii_195.pdf (accessed: 2017–03–21).

Statistics Norway (1937a): *Sundhetstilstanden og medisinalforholdene 1935*. Oslo. www.ssb.no/a/histstat/nos/nos_ix_125.pdf (accessed: 2017–03–21).

Statistics Norway (1937b): *Norges civile, geistlige, rettslige og militære inndeling 1. juli 1937*. Oslo. www.ssb.no/a/histstat/nos/nos_ix_119.pdf (accessed: 2017–03–21).

Statistics Norway (1939a): *Sundhetstilstanden og medisinalforholdene 1936*. Oslo. www.ssb.no/a/histstat/nos/nos_ix_153.pdf (accessed: 2017–03–21).

Statistics Norway (1939b): *Sunnhetstilstanden og medisinalforholdene 1937*. Oslo. www.ssb.no/a/histstat/nos/nos_ix_178.pdf (accessed: 2017–03–21).

Statistics Norway (1940): *Sunnhetstilstanden og medisinalforholdene 1938*. Oslo. www.ssb.no/a/histstat/nos/nos_ix_194.pdf (accessed: 2017–03–21).

Statistics Norway (1941a): *Sunnhetstilstanden og medisinalforholdene 1939*. Oslo. www.ssb.no/a/histstat/nos/nos_x_021.pdf (accessed: 2017–03–21).

Statistics Norway (1941b): *Norges sivile, geistlige og rettslige inndeling 1. juli 1941*. Oslo. www.ssb.no/a/histstat/nos/nos_x_020.pdf (accessed: 2017–03–21).

Statistics Norway (1943): *Sunnhetstilstanden og medisinalforholdene 1940*. Oslo. www.ssb.no/a/histstat/nos/nos_x_055.pdf (accessed: 2017–03–23).

Statistics Norway (1946): *Sunnhetstilstanden og medisinalforholdene 1941*. Oslo. www.ssb.no/a/histstat/nos/nos_x_097.pdf (accessed: 2017–03–23).

Statistics Norway (1948): *Sunnhetstilstanden og medisinalforholdene. Hovedresultatene 1942-1944*. Oslo. www.ssb.no/a/histstat/nos/nos_x_160.pdf (accessed: 2017–03–20).

Statistics Norway (1949): *Sunnhetstilstanden og medisinalforholdene 1945. Og sammendragstabeller 1941-1945*. Oslo. www.ssb.no/a/histstat/nos/nos_x_179.pdf (accessed: 2017–03–23).

Statistics Norway (1951): *Norges sivile, geistlige, rettslige og militære inndeling 1. januar 1951*. Oslo. www.ssb.no/a/histstat/nos/nos_xi_049.pdf (accessed: 2017–03–21).

Statistics Norway (2015): *Tabell: 10759: Befolkningsendringer, etter kjønn og bygder/byer (K) (avslutta serie)*. www.ssb.no/statistikkbanken/selectvarval/Define.asp?SubjectCode=al&ProductId=al&MainTable=BefKomTabPub&SubTable=Kommun1&PLanguage=0&Qid=0&nvl=True&mt=1&pm=&gruppe1=Hele&gruppe2=Hele&gruppe3=Hele&aggreg1=&aggreg2=&aggreg3=&VS1=Kommuner1906&VS2=Kjonn3&VS3=BygderByer&CMSSubjectArea=befolkning&KortNavnWeb=folkemengde&StatVariant=&TabStrip=Select&checked=true (accessed: 2017–04–05).

Statistics Norway (2017a): *Tabell: 06913: Folkemengde 1. januar og endringer i kalenderåret (K)*. www.ssb.no/statistikkbanken/selectvarval/Define.asp?SubjectCode=al&ProductId=al&MainTable=Folkemengd1951&SubTable=Kommun1&PLanguage=0&Qid=0&nvl=True&mt=1&pm=&gruppe1=KommNyHist&

aggreg1 = NO & VS1 = Kommuner1951 & CMSSubjectArea = befolkning & KortNavnWeb=folkemengde&StatVariant=&TabStrip=Select&checked=true (accessed: 2017–04–05).

Statistics Norway (2017b): *Kopi av kommuneendringer_1838_2017*. Oslo/Kongsvinger. www.ssb.no/offentlig-sektor/kommunekatalog/endringer-i-de-regionale-inndelingene/_attachment/291670?_ts=159b175e638 (accessed: 2017–02–20).

Store norske leksikon (2015): *Norsk utdanningshistorie*. www.snl.no/Norsk_utdanningshistorie (accessed: 2017–05–11).

NIPH – The Norwegian Institute of Public Health (2015): *Flere bør ta influensavaksinen*. www.fhi.no/nyheter/2015/flere-bor-ta-influensavaksinen-/ (accessed: 2017–06–13).

NIPH – The Norwegian Institute of Public Health (2016): *Influenza – Fact sheet about seasonal influenza*. www.fhi.no/en/id/influensa/seasonal-influenza/influenza---fact-sheet-about-season/ (accessed: 2017–02–06).

The Norwegian Mapping Authority (2017): *Fylkes- og kommuneoversikt*. www.kartverket.no/kunnskap/fakta-om-norge/Fylker-og-kommuner/Tabell/ (accessed: 2017–03–24).

Wooldridge, J. M. (2012): *Introductory Econometrics. A Modern Approach*. 5th ed. Mason, OH: Cengage Learning.

WHO – World Health Organization (2016): *Influenza (Seasonal)*. www.who.int/mediacentre/factsheets/fs211/en/ (accessed: 2017–06–13).

Wüst, M. (2012): "Early interventions and infant health: Evidence from the Danish home visiting program". *Labour Economics* 19(4), 484–495.

Appendix

A1 Tables

Table A1: Summary Statistics: County-Specific Influenza Morbidity

	Observations	Mean	Std. dev.	Min	Max
Akershus	1,980	1.881	3.171	0.000	29.307
Aust-Agder	1,716	1.365	3.658	0.000	76.293
Bergen	132	1.093	3.783	0.000	40.687
Buskerud	1,980	2.322	4.460	0.000	52.813
Finnmark	1,848	0.980	3.041	0.000	62.252
Hedmark	3,036	1.342	3.379	0.000	93.110
Hordaland	3,300	1.339	9.191	0.000	330.530
Møre og Romsdal	4,092	1.057	2.216	0.000	33.353
Nord-Trøndelag	3,036	1.565	4.183	0.000	120.393
Nordland	5,940	1.483	3.572	0.000	106.100
Oppland	3,432	1.392	2.785	0.000	31.826
Oslo	132	0.889	1.698	0.000	12.835
Rogaland	2,112	1.495	2.920	0.000	44.044
Sogn og Fjordane	3,168	1.155	2.706	0.000	53.150
Sør-Trøndelag	2,904	1.008	2.496	0.000	38.187
Telemark	2,376	1.921	3.514	0.000	69.041
Troms	2,376	1.170	2.874	0.000	62.781
Vest-Agder	1,848	1.167	3.071	0.000	41.955
Vestfold	1,584	2.274	5.598	0.000	118.320
Østfold	2,112	1.695	3.601	0.000	45.285

Notes: This table presents the number of monthly influenza observations per 1,000 inhabitants, the mean value, the standard deviation, and the minimum and maximum values of influenza prevalence in each Norwegian county in the period 1935-1945. The sample includes all non-missing observations for children born between 1935 and 1945 in municipalities that got a mother and child health care center before 1956.

Table A2: Selection Effects over the Life Cycle

	Stillbirths		Infant mortality		Long-run outcome missing	
	(1)	(2)	(3)	(4)	(5)	(6)
	Males	Females	Males	Females	Education	Log Earnings
Flu	-0.735	-0.034	3.945	4.423	0.000	0.000
	(1.470)	(0.858)	(2.080)	(2.697)	(0.000)	(0.000)
Observations	377	514	8,952	8,952	47,666	47,666
Adjusted R^2	0.880	0.883	0.540	0.528	1.000	1.000

Notes: Each column represents a different regression of the outcome indicated by the column name on a measure of influenza exposure in the in utero period. For stillbirths and infant mortality, influenza exposure is measured as the average influenza prevalence per 1,000 inhabitants the infants born the given year were exposed to in utero . All specifications include dummies for the year and medical district of birth. For missing long-run outcomes, influenza exposure is measured as the prevalence per 1,000 inhabitants during each trimester. Only the estimated impact of second trimester exposure to maternal influenza is reported. All specifications include gender, birth order, and dummies indicating whether the mother was unmarried and whether the parents' education levels and age at the time of the child's birth were below average, as well as the ratios of doctors to inhabitants at the time of the child's birth and students to teachers in the year the child was enrolled at school. Additionally, I control for cohort, district and mother fixed effects, as well as a district-specific linear time trend. Effects are therefore identified as deviations from the trend. In order to estimate the effect of influenza rates deviating from the average seasonality in the district, I include month-district fixed effects. The sample includes all non-missing observations for children born between 1935 and 1945 in municipalities that got a mother and child health care center before 1956. Robust standard errors clustered at the level of the medical district of birth are shown in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A3: Earnings Measures in Levels

	Earnings (in levels)		
	(1)	(2)	(3)
	1967-2010	Age 31-40	Age 41-50
Centers*Flu trimester 2	183.142	411.943	-1026.442
	(928.593)	(1045.337)	(1373.454)
Observations	47,666	46,947	47,666
Adjusted R^2	0.450	0.602	0.413

Notes: Each column represents a different regression of the variable indicated by the column name on the influenza prevalence per 1,000 inhabitants during each trimester, a dummy indicating access to a mother and child health care center, and an interaction between each trimester and the center dummy. Only the coefficient on the interaction effect for the second trimester is reported. All specifications include gender, birth order, and dummies indicating whether the mother was unmarried and whether the parents' education levels and age at the time of the child's birth were below average, as well as the ratios of doctors to inhabitants at the time of the child's birth and students to teachers in the year the child was enrolled at school. Additionally, I control for cohort, district and mother fixed effects, as well as a district-specific linear time trend. Effects are therefore identified as deviations from the trend. In order to estimate the effect of influenza rates deviating from the average seasonality in the district, I include month-district fixed effects. The sample includes all non-missing observations for children born between 1935 and 1945 in municipalities that got a mother and child health care center before 1956. Robust standard errors clustered at the level of the medical district of birth are shown in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A4: Influenza as a Single Shock

	Education (margin)		Log Earnings		
	(1)	(2)	(3)	(4)	(5)
	Intensive	Extensive	1967-2010	Age 31-40	Age 41-50
Flu trimester 2	0.015	0.003	0.002	-0.011	-0.011
	(0.020)	(0.002)	(0.008)	(0.008)	(0.007)
Observations	38,866	38,866	47,047	44,927	45,549
Adjusted R^2	0.285	0.296	0.381	0.485	0.345

Notes: 'Intensive margin' refers to completed years of schooling. 'Extensive margin' refers to whether individuals took more than the mandatory level of schooling. Each column represents a different regression of the outcome indicated by the column name on the influenza prevalence per 1,000 inhabitants during each trimester. Only the estimated impact of influenza in the second trimester is reported. All specifications include gender, birth order, and dummies indicating whether the mother was unmarried and whether the parents' education levels and age at the time of the child's birth were below average, as well as the ratios of doctors to inhabitants at the time of the child's birth and students to teachers in the year the child was enrolled at school. Additionally, I control for cohort, district and mother fixed effects, as well as a district-specific linear time trend. Effects are therefore identified as deviations from the trend. In order to estimate the effect of influenza rates deviating from the average seasonality in the district, I include month-district fixed effects. The sample includes all non-missing observations for children born between 1935 and 1945 in municipalities that got a mother and child health care center before 1956. Robust standard errors clustered at the level of the medical district of birth are shown in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A5: Mother and Child Health Care Centers as a Single Shock

	Education (margin)		Log Earnings		
	(1)	(2)	(3)	(4)	(5)
	Intensive	Extensive	1967-2010	Age 31-40	Age 41-50
Centers	0.342*	0.044	-0.023	-0.106	-0.077
	(0.144)	(0.025)	(0.055)	(0.072)	(0.072)
Observations	38,866	38,866	48,637	44,927	47,099
Adjusted R^2	0.285	0.296	0.367	0.485	0.331

Notes: 'Intensive margin' refers to completed years of schooling. 'Extensive margin' refers to whether individuals took more than the mandatory level of schooling. Each column represents a different regression of the dependent variable indicated by the column name on a dummy variable indicating whether the medical district where the infant was born had access to a mother and child health care center before the infant became one year old. All specifications include gender, birth order, and dummies indicating whether the mother was unmarried and whether the parents' education levels and age at the time of the child's birth were below average, as well as the ratios of doctors to inhabitants at the time of the child's birth and students to teachers in the year the child was enrolled at school. Additionally, I control for cohort, district and mother fixed effects, as well as a district-specific linear time trend. Effects are therefore identified as deviations from the trend. In order to estimate the effect of influenza rates deviating from the average seasonality in the district, I include month-district fixed effects. The sample includes all non-missing observations for children born between 1935 and 1945 in municipalities that got a mother and child health care center before 1956. Robust standard errors clustered at the level of the medical district of birth are shown in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table A6: Removing Doctors-to-Inhabitants from the Control Variables

	Education (margin)		Log Earnings		
	(1)	(2)	(3)	(4)	(5)
	Intensive	Extensive	1967-2010	Age 31-40	Age 41-50
Centers	0.373*	0.034	-0.030	-0.117	-0.084
	(0.162)	(0.028)	(0.071)	(0.081)	(0.098)
Flu trimester 1	0.062*	0.001	0.013	0.004	-0.004
	(0.027)	(0.004)	(0.009)	(0.015)	(0.009)
Flu trimester 2	0.010	0.001	0.013	-0.014	-0.009
	(0.025)	(0.004)	(0.011)	(0.012)	(0.012)
Flu trimester 3	-0.004	-0.006	-0.015	-0.005	-0.004
	(0.033)	(0.006)	(0.010)	(0.013)	(0.010)
Centers*Flu trimester 1	-0.038	-0.001	0.004	0.001	0.016
	(0.034)	(0.004)	(0.012)	(0.018)	(0.011)
Centers*Flu trimester 2	0.010	0.003	-0.020	0.004	-0.002
	(0.034)	(0.005)	(0.013)	(0.013)	(0.013)
Centers*Flu trimester 3	0.004	0.004	0.015	0.001	-0.003
	(0.033)	(0.007)	(0.011)	(0.016)	(0.011)
Observations	38,866	38,866	47,047	44,927	45,549
Adjusted R^2	0.286	0.296	0.381	0.485	0.346

Notes: 'Intensive margin' refers to completed years of schooling. 'Extensive margin' refers to whether individuals took more than the mandatory level of schooling. Each column represents a different regression of the variable indicated by the column name on the influenza prevalence per 1,000 inhabitants during each trimester, a dummy indicating access to a mother and child health care center, and an interaction between each trimester and the center dummy. All specifications include gender, birth order, and dummies indicating whether the mother was unmarried and whether the parents' education levels and age at the time of the child's birth were below average, as well as the ratio of students to teachers in the year the child was enrolled at school. Additionally, I control for cohort, district and mother fixed effects, as well as a district-specific time trend, which is specified as linear unless the row name indicates otherwise. Effects are therefore identified as deviations from the trend. In order to estimate the effect of influenza rates deviating from the average seasonality in the district, I include month-district fixed effects. The sample includes all non-missing observations for children born between 1935 and 1945 in municipalities that got a mother and child health care center before 1956. Robust standard errors clustered at the level of the medical district of birth are shown in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

A2 Municipalities in 1930 and 1980

When deciding how to treat municipality changes between 1934 and 1980, I have relied on the descriptions of each change provided by Juvkam (1999). As a point of departure, I used the Excel file provided by Statistics Norway (2017b), which builds on Juvkam (1999). However, the file contained a rather large amount of errors. I therefore corrected the errors I came across, verifying the corrections using Juvkam (1999), the overview of current municipality numbers and names provided by The Norwegian Mapping Authority (2017), as well as other changes in the Excel file. I have attempted to ensure that all changes seem plausible after the correction.

I disregard all small border regulations, as well as all municipality mergers that are later undone. In cases of mergers between municipalities from different medical districts, I consider the resulting municipality as part of the medical district of the largest of the merged municipalities. In cases where a municipality splits into multiple parts, I follow only the largest part, unless the smaller part becomes an entirely new municipality, in which case I note which medical district the new municipality stems from.