Long run labor market outcomes of a childhood health shock: high socioeconomic status parents can mitigate the impact

Tine M. Eriksen, Amanda P. Gaulke, Jannet Svensson, Niels Skipper, & Peter R. Thingholm

Draft: May 2022

Abstract

While a vast literature examines the impact of childhood health shocks, questions remain about the distribution of effects and the role of parents in mitigating or exacerbating the impact. Studying this is challenging as poor health is non-random. Our paper overcomes this challenge by using the quasi-random onset of Type 1 Diabetes (T1D) in childhood. Using Danish administrative registry data, we first use OLS to find differences, on average, in adult employment and labor market income by matching people with and without T1D on exact date of birth and sex. Second, results from the causal forest analyses show considerable heterogeneity. For example, we find the conditional average treatment effects for labor income range from DKK -68,122 to DKK 10,519. In a setting with universal and equal access to health care, we document a clear and economically meaningful socioeconomic gradient in effects, as high socioeconomic status parents mitigate the impact.

1. Introduction

Health is likely an important mechanism in driving the intergenerational transmission of inequality (Bowles and Gintis, 2002). Case, Lubotsky, and Paxson (2002) show that the income gradient for health previously documented among adults also exists in early childhood. Additionally, they find that lower socioeconomic status children become even more penalized over time. Currie and Stabile (2003) test for potential mechanisms and determine that health shocks do not impact high and low socioeconomic status children differently, but rather the previous relationship seems to be driven by lower socioeconomic status children experiencing more negative health shocks.

Researchers have documented that a wide variety of shocks that occur in utero or in early childhood can impact wages and employment in adulthood (see Almond, Currie, and Duque (2018) for an overview). While most of the literature finds negative impacts from childhood health shocks, Gensowski, Nielsen, Nielsen, Rossin-Slater and Wüst (2019) find positive impacts from the 1952 polio epidemic in Denmark. Specifically, they find positive effects on completing a university degree and working in a white-collar or computer intensive job.

Despite the well-established health income gradient, the vast research on the long run impacts of childhood health shocks has largely focused on average treatment effects and thus could be missing important treatment effect heterogeneity, for instance by paternal and maternal socioeconomic status. In this paper we ask, can high socioeconomic status parents mitigate effects from a childhood health shock that has been previously found to have negative average treatment effects? We use childhood onset of Type 1 Diabetes (T1D), rich administrative registry data from Denmark, and multiple methodologies to answer this question.

Our first contribution is to document considerable heterogeneity in effects that are missed by solely focusing on the average treatment effect. We do this by estimating average

treatment effects using OLS regressions and the full distribution of conditional average treatment effects using the causal forest methodology (Wager and Athey, 2018). Specifically, we find that the OLS results for labor market income suggest that the average treatment effect is a decline of DKK 32,422 while the causal forest shows the full distribution of conditional average treatment effects varies from DKK -68,122 to DKK 10,519. When the outcome is a dummy variable for employment (having a positive labor income), the OLS estimate is an 8.7 percent decline, but the causal forest shows the full distribution of the conditional average treatment effects ranges from a 22.7 percent decline to no impact. Thus, the previously found significant and negative labor market impacts of a childhood T1D diagnosis (for example, see Persson et al. (2016)) masks considerable heterogeneity in effects.

Not only do we document important heterogeneity in conditional average treatment effects, but we find evidence of some people not being negatively impacted. This means that some parents can mitigate the impact of a childhood health shock, even when the average treatment effect is negative. If we could pinpoint which subgroups of parents are mitigating the impact, we could leverage their strategies to reduce inequality stemming from childhood health shocks.

Motivated by the income health gradient and Almond, Currie, and Duque (2018) which states "a greater understanding of the way that shocks and disadvantage interact, and of the role of parents in responding to them, is highly desirable", we test for differences in relative effects by parental socioeconomic status. Thus, our second contribution is to show there is a socioeconomic gradient in relative effects by leveraging data on both mothers and fathers. The pattern of relative effects is similar for both maternal and paternal income quartiles in that having a parent who is in a higher income quartile leads to smaller penalties in labor market outcomes. Having a more educated father generally leads to smaller penalties in labor market outcomes (except for vocational education), whereas the pattern for mothers is an

inverse U-shape. Having a mother with a master's degree or above results in a larger penalty than having a mother with a bachelor's degree. Thus, we find clear evidence of a socioeconomic gradient in long run impacts among people with childhood onset of T1D, even in a country with universal access to healthcare. This contrasts with Currie and Stabile (2003) which did not find evidence of health shocks differentially impacting children by socioeconomic status.

Related literature has shown variation in parental responses to other health shocks (for example, Datar et al. (2010), Hsin (2012) and Restrepo (2016) focus on low-birthweight and Guo and Zhang (2021) define poor health as having any of the following conditions by age 18: migraine, rash, disability, serious hearing difficulties, heart disease, pollen allergy, neurasthenia, hypertension, or alcoholism).[1] These papers importantly show variation in parental responses using direct measures of parental investments. We instead focus on the heterogeneity in relative effects on labor market outcomes by maternal and paternal characteristics and show there is a socioeconomic gradient in that.

We investigate three key mechanisms that could explain impacts on labor market outcomes. First, despite universal access to health care and much lower out-of-pocket costs associated with treatment than in the United States, we find suggestive evidence that parental socioeconomic characteristics impact adult health. While the 95% CIs overlap, the point estimates suggest that individuals with fathers or mothers in lower income quartiles have worse T1D related health outcomes as adults. Specifically, they have higher HbA1c levels (worse glucose control) and are more likely to be hospitalized due to diabetic coma or diabetic ketoacidosis. Health capital likely matters because previous work has found worse

[1] Grätzand and Torche (2016) instead focus on how parental investments interact with child ability and Houmark, Ronda and Rosholm (2020) focus on how parents investment interacts with genetics. Literature has also documented important heterogeneity by family background for non-health related shocks (for example, see Havnes and Mogstad (2011), Løken, Mogstad, and Wiswall (2012), Carneiro, Løken, and Salvanes (2015) and Havnes and Mogstad (2015)).

educational outcomes for children and adolescents with more poorly controlled T1D (for example, see Skipper et al. (2019), Eriksen et al. (2020), and Lindkvist et al. (2021)). The second mechanism we test for is therefore human capital accumulation. We find clear evidence of those with the smallest impact on labor market outcomes having the highest levels of educational attainment. Lastly, we test for differences in impacts by T1D duration, and find that those diagnosed at younger ages (less than 9 years old) have larger negative impacts.

The rest of the paper is structured as follows. Section 2 provides background on T1D and the institutional setting in Denmark. Section 3 describes the Danish administrative registry data and provides some descriptive statistics. Section 4 discusses the OLS regression (homogenous effects) and the causal forest methodology (heterogeneous effects). Section 5 discusses the results from both the OLS regressions and causal forests, along with evidence of a clear socioeconomic gradient in effects. Section 6 discusses the mechanisms behind the variation found in the causal forest analyses, including health outcomes, educational outcomes and T1D duration. Section 7 summarizes and concludes.

2. Background

2.1 Background on T1D

T1D[2] is a chronic health condition in which the immune system kills the insulin producing cells in the pancreas. Insulin is a metabolic hormone needed to allow sugar (glucose) to enter the cells in the body to produce energy. In healthy individuals, the glucose concentration in the bloodstream is constantly kept within a narrow interval. This is done by secreting insulin when glucose levels rise, i.e., when eating and drinking (primarily carbohydrates), to get

[2] T1D has previously been known as juvenile diabetes or insulin dependent diabetes.

glucose levels down, and by releasing sugar (primarily from the liver) outside of meals (or when fasting). In individuals with T1D, the endogenous insulin supply is absent.

Managing T1D is time consuming and complex. As there is no cure, the aim is to keep glucose levels as close to normal as possible by injecting insulin to counter the rises in glucose from both eating and the endogenous release of sugar. The amount of insulin needed depends on the food intake[3] and amount of exercise, just to name a few things. In addition, numerous small blood samples from finger pricks or a continuous glucose monitor must be used to check glucose levels and adjust insulin dosages, as necessary.

Even with insulin treatment, individuals with T1D have chronically elevated glucose levels (known as hyperglycemia) compared with individuals without T1D. Hyperglycemia is strongly associated with so-called late complications such as kidney disease/failure, blindness, amputation, cardiovascular disease, etc. (DCCT, 1993). At the other end of the spectrum, administering too much insulin will cause a low blood sugar (known as hypoglycemia), which can lead to seizures and coma. Thus, glucose levels need to be checked frequently, including during the night. Hypoglycemia is treated by ingesting a sugary snack or beverage. This constant alertness required to keep the blood sugar in check has been shown to negatively affect parental sleep (Cobry and Jaser, 2019, and Pillar et al., 2003), and it obviously requires mental resources to keep the condition well-managed.

T1D is – in contrast to the more common type 2 diabetes – not related to lifestyle or sedentary behavior. It most often presents in childhood, and it is, second to asthma, the most common chronic physical health condition in children and adolescents in most of the Western

[3] The amount of insulin needed is typically calculated as an insulin-to-carbohydrate ratio, i.e., a specific amount of insulin needs to go with a certain amount of carbohydrates. However, this ratio is different from individual to individual, and even varies over the time of day. The type of carbohydrate and the protein and fat content of the food has implications for the correct insulin dosage.

world[4]. The exact cause of the disease is unknown. While there is ongoing research to determine the origin, Regnell and Lernmark (2017) note, "The aetiology of beta cell autoimmunity is still unclear". Genetics play a smaller role than for other chronic health conditions as Pociot and Lernmark (2016) note there is only a three percent risk of T1D in children of mothers with the condition. The condition is characterized by a rapid onset in children and adolescents and affected individuals will not go undiagnosed. Before insulin treatment was discovered in the early 20th century, a patient newly diagnosed with T1D had an average life expectancy of around 2 years (Hakim et al., 2013).

To illustrate that the onset of T1D is sudden, we show results from Thingholm et al. (2020) which uses an event-study design to compare school absenteeism of children eventually diagnosed with T1D to matched comparison children (who are then assigned a pseudo diagnosis date) based on sex and date of birth. The graph is informative in several ways that are key to our identifying assumption of T1D constituting a near-random health shock. It illustrates that the child's symptoms (severe enough to affect school absenteeism) are only present in the months very close to diagnosis (four-months prior is the first month with a significant difference). Related, there seem to be no differences in the underlying health of children who develop T1D, based on their pre-diagnosis levels of school absenteeism (i.e., more than four months before onset). Further, Figure 1 shows that after the onset of diabetes, these children have systematically and significantly more school absenteeism (roughly 50% more than the comparison children). This points to the fact that something is now different for these children.

As we will show later, there is no SES gradient in T1D as our F-tests rule out differences between the treatment and comparison groups. Eriksen, Gaulke, Skipper, and Svensson

[4] In the US, around 1.25 million children and adults live with type 1 diabetes, with an estimated annual cost of US$ 14.4 billion (Tao et al. (2010)). In Denmark, approximately 32,000 people are diagnosed with type 1 diabetes, including 3,500 children.

(2021) find that mothers and fathers of children who will be diagnosed with T1D had comparable wage earnings to parents of sex and age-matched children prior to the diagnosis. Further, Eriksen et al. (2021) shows that the children who are diagnosed with T1D have no significant differences in 5-minute APGAR scores compared with children who will not be diagnosed and both groups are equally likely to be of low birth weight. Lastly, later in the paper we provide more evidence consistent with no differences in underlying health prior to diagnosis. There are no significant differences between our treated and comparison children in terms of admission to the hospital, visits to the general practitioner (primary care physician) and the probability of pharmacy claims prior to diagnosis.

2.2 Background on the Institutional Setting

In Denmark, the financial cost of managing T1D is low, especially compared with the costs faced in the United States. There is no cost for in-patient or out-patient care and all medical devices are provided free of charge. However, medicine is not free and thus insulin has an associated out-of-pocket cost. The median yearly out-of-pocket expenditure for insulin in 2016 was approximately $229 and low-income families can apply to have the cost waived.

In Denmark extensive legislation exists to ensure equal rights both in- and outside of the labor market concerning aspects related to gender, health, age, etc. Workers are not obliged to disclose the diagnosis to their employer, and in general, discrimination based on health conditions is illegal. This applies to both formal and informal discrimination; however, the latter can be hard to verify.

3 Data and descriptive statistics

The data are from several Danish registries. All children diagnosed with T1D in Denmark since 1996 are included in a national registry called DanDiabKids. It contains information on all Danish children and adolescents diagnosed with T1D who are seen at pediatric

endocrinology clinics. Transfer to an adult endocrinology clinic usually happens at age 18, although some transfer at age 16. Each child is identified through a person-specific ID (equivalent to a social security number). With these data, we observe the exact date of diagnosis along with clinical characteristics of the children and adolescents collected at annual follow-up visits in pediatric endocrinology clinics (ambulatories). We complement these data with hospitalization records which start in 1976 (*Landspatientregisteret*). This allows us to observe the exact day of the contact as well as the diagnosis associated with the visit (ICD-10 codes and before that ICD-8 codes). These two data sources are used to identify the children diagnosed with T1D. To ensure that the children are old enough to examine long-run labor market outcomes, we focus on the cohorts of children born from 1977 to 1987.

Through population registries within Statistics Denmark, the diagnosis data is augmented with information on demographic and socioeconomic characteristics. Further, children can be linked to their parents. We observe age, sex, educational attainment, income, and immigrant or descendant status for all individuals. For the children, we further observe birth order and the number of siblings. All data are recorded at the yearly level. We restrict the sample to include only native Danes to ensure that we have 1) accurate information about the child's medical history as we can only observe health records for care provided in Denmark and 2) information on parental background characteristics prior to the child's birth so we can compare our treatment and comparison group means without worrying that characteristics are being impacted by the T1D diagnosis. This is important because Eriksen et al. (2021) find impacts on parental labor supply and mental health because of the child's diagnosis. Using income measured in the year prior to the child's birth does result in some missing data (one third have missing data) since we can only go back to 1980 for income data and some cohorts are born before that. We have, however, tested for whether there are differences in missing income data across our treatment and comparison parents and find no statistical difference.

As the information on other background characteristics is only available from 1986 and onwards, we limit our data to children and adolescents who are diagnosed from 1986 to 2004. We identify 1,810 children who are diagnosed with T1D before they turn 18. The mean age at onset is 11.8 years (SD 3.8). For each child with T1D, we identify and match five comparison children (9,050 individuals) who are not diagnosed with T1D before their 30th birthday. The matching is performed on the exact date of birth and sex.

To lend further credibility to the assumption that T1D is conditionally random we use an event study analysis to compare healthcare utilization among individuals with and without a T1D diagnosis, each month from two years prior to diagnosis through two years after diagnosis, see Figure 2. Panel A shows there are no differences in the probability of admission to a hospital prior to diagnosis, Panel B shows there is no difference in visits to the general practitioner prior to diagnosis, and Panel C shows there are no differences in pharmacy claims prior to diagnosis. In all three cases, children who are diagnosed are significantly more likely to utilize that type of health care in the month of diagnosis and the months after diagnosis. These results further lend credibility to the claim that there were no differences in underlying health for children who are diagnosed with T1D.

Lastly, we compare observable characteristics of the children and parents between the treatment and comparison groups; see Table 1. The two groups seem remarkably similar. We regress an indicator variable for T1D onset on the characteristics listed in Table 1 to see if we can predict which children will have childhood onset of T1D. Results are shown in Table 2. While there are some significant coefficients, this is to be expected given the number of coefficients we are estimating, and the joint F-test suggests no overall differences between the two groups.

4 Empirical Strategy

In this section we present our empirical strategy to estimate the impact of a childhood health shock on adult labor market outcomes. Ultimately, we rely on the conditional randomness of a T1D diagnosis and employ a causal forest methodology, which is a novel non-parametric approach that estimates heterogeneous treatment effects. A causal forest is well suited for our setting, as we have a relatively large number of covariates, and no a priori theoretical ranking of their importance to the outcome.

Investigating whether family background matters for how an individual is impacted by adverse health events in childhood is not straightforward. The ideal experiment from a statistical sense would be to randomly allocate adverse events to the general population and observe the differences in earnings at (e.g.) age 30. However, we obviously cannot do that. Our observed data consists of (X_i, Y_i, D_i), where X_i are observed characteristics, Y_i is observed outcome and D_i is an indicator equal to 1 if individual *i* experiences a health shock in childhood. Everyone has two potential outcomes $Y_i^D, D \in (0,1)$. In our setting, it is the outcome for a particular individual if she experienced a health shock ($D_i = 1$), or if she has not experienced the health shock ($D_i = 0$). The treatment effect we seek to estimate is defined as

$$\tau = E[Y_i^1 - Y_i^0]$$

Since one of our research questions is the extent to which this treatment effect varies with parental background, we need to allow the treatment effect to be a function of our set of parental demographic covariates, X_i:

$$\tau(x) = E[Y_i^1 - Y_i^0 \,|X_i = \mathrm{x}]$$

The challenge is that we only observe either Y_i^1 or Y_i^0 for each individual *i*, and without further assumptions on the relationship between D_i and Y_i^D we cannot estimate the treatment effect, $\tau(x)$. Hence, to estimate the impact of a health shock, one needs a health shock that is

independent of the potential outcomes - $(Y_i^1, Y_i^0) \perp D_i$. In this paper we leverage such an independence between T1D and potential outcomes.

Estimation

Under the (conditional)[5] independence assumption we can estimate τ, the homogenous impact of diabetes on labor market outcomes at age 30 in a linear model using ordinary least squares

$$Y_i = \tau D_i + \boldsymbol{X}_i\beta + \varepsilon_i \qquad (1)$$

where Y_i is one of the two outcomes of interest measured at age 30 (labor income or reporting any positive labor income). D_i is an indicator for being diagnosed with T1D in childhood (before age 18). $\boldsymbol{X}_i$ is a set of parental demographic covariates, and ε_i is a normally distributed error term.

Several approaches have traditionally been used to estimate heterogeneous treatment effects. First, heterogeneous treatment effects can be estimated by separately estimating (1) for each value of $\boldsymbol{X}_i$:

$$Y_i = \tau(x)D_i + \varepsilon_i \; if \; \boldsymbol{X}_i = x \qquad (2)$$

While this is a very flexible approach, as it does not impose any additional structure on the relationship between Y_i and D_i, there are several issues with this approach. First, increasing the dimension of $\boldsymbol{X}_i$ can lead to regions with very sparse density, and for continuous X_i's perfectly subsetting is impossible. Hence continuous variables are often discretized. Second, there exist no ex-ante guidance on how to subset the data, how to discretize the continuous variables, and a researcher ex-ante must motivate the split and ex-post correct for multiple hypothesis testing.

Alternatively, to estimate, $\tau(x)$, one can estimate a linear model like (1) where D_i is fully interacted with the set of parental demographics. This again requires the researcher to

[5] We could equivalently impose a conditional independence assumption.

arbitrarily discretize continuous variables. In addition, as the dimension of $\boldsymbol{X}_i$ increases, the estimation by OLS becomes infeasible.

CART and Causal Forests

In this paper we choose a different approach and estimate the $\tau(x)$ by causal forests. As described by Wager and Athey (2018) both trees and forests can be thought of as a k-nearest neighbor method, where the closeness is defined in the context of a decision tree. Here the observations are divided into subsets or *leaves* using data-driven sample splits. In contrast to traditional methods, where a prespecified metric such as Euclidean distance is used, the closest observation to a given point is the observations that end up in the same leaf. Wager and Athey (2018) also note that causal forests are a way to increase power.

To understand how the procedure works, it is helpful to recall how a standard *classification and regression tree* (CART) is used for prediction. A CART stratifies the set of characteristics into a number of regions often referred to as *leaves*. The prediction for a given observation is typically the mean of the outcome in the region to which the splits assign it. Let $L(X_i)$ denote end leaf of individual i, with characteristics X_i. The prediction of Y_i can then be written as

$$\hat{\mu}(X_i) = \sum_{i=1}^{N} \alpha_i(X_i)\ Y_i \qquad (3)$$

where

$$\alpha_i(X_i) = \frac{1[X_i \in L(X_i)]}{|i{:}X_i \in L(X_i)|}, \qquad (4)$$

or the share individual i constitutes of the total number of individuals in leaf $L(X_i)$. In other words, the prediction from a CART is a local average of all the $Y_i's$ that are "close" (in the same leaf as) to individual i.[6] The leaves are obtained by recursively splitting the feature-space, and at each decision node the object is to maximize heterogeneity of "child-nodes".

[6] Conceptually the CART can be thought of as a non-parametric regression with a decision tree as the distance kernel

Several procedures exist for how to place the splits (see, for example, Hastie, Tibshirani, and Friedman (2009) for details).

Causal trees extend this idea to estimation of heterogenous treatment effects. That is, as an estimator of heterogenous treatment effects, the causal tree calculates within each leaf the mean outcome for those who are treated, and subtracts the mean outcome for those who are not treated:

$$\hat{\tau}(X_i = x) = \frac{1}{|i:D_i=1, X_i \in L(X_i)|} \sum_{i=1}^{N} 1[D_i = 1, X_i \in L(X_i)] Y_i - \frac{1}{|i:D_i=0, X_i \in L(X_i)|} \sum_{i=1}^{N} 1[D_i = 0, X_i \in L(X_i)] Y_i \quad (5)$$

At each node, the Causal Tree maximizes the variation in eq. (5). Wager and Athey (2018) further lay out the additional properties that need to hold. It is critical for our application that we ensure the presence of treated and comparisons in each leaf – normally known as *overlap* – and this is ensured by imposing the rule that at least a fraction of $\alpha \in (0,1)$ of the individuals in each leaf is either treated or part of the comparison group.[7]

An obvious question is how to determine the "best" tree. *Causal Forests* build on the insights from Breiman (2001) that averaging over several trees is an improvement on estimating a single tree. Hence, after estimating B causal trees, and obtaining an estimate for each of these $\hat{\tau}_b(x)$, the Causal Forest aggregate these by averaging them to obtain an estimator of the individual conditional treatment effect:

$$\hat{\tau}(x) = B^{-1} \sum_{b=1}^{B} \hat{\tau}_b(x)$$

Wager and Athey (2018) shows pointwise consistency, and that the predictions made by a causal forest are asymptotically normal and unbiased. As previously stated, this is the fundamental reason for why we apply this technique, as it implies that we can do post

[7] This is known as α-regularization. Another important distinction from standard CART estimation is the use of honest trees – also called *double sample* trees where the training data is split into two groups; the first used to determine the splits, the second used to evaluate the treatment effect. For details see Wager and Athey (2018).

estimation analysis without being concerned about multiple hypothesis testing. In practice we implement the causal forest with honest trees (Athey and Imbens, 2016), grow 2,000 trees and impose that at a minimum 5% of individuals in any leaf must be either treated or non-treated (α =0.05). The final part of our analysis is to investigate variation in treatment effects by comparing mean estimated conditional treatment effects (CATEs) across subgroups.[8]

A concern when estimating CATEs is that even in the absence of true treatment effect heterogeneity, the CATEs will be estimated with some distribution (noise). To assess if this is the case, we divide our sample into quartiles based on the predicted treatment effect. For each quartile, we then estimate the group average treatment effect (GATE) and compare them across quartiles. If these are different, this is evidence of treatment effect heterogeneity.

5 Results

5.1 Main Results

We start out by presenting evidence on the homogeneous treatment effects (OLS) of being diagnosed with T1D during childhood. Table 3 shows the results for our two main outcomes; labor income (Panel A) and having positive labor income (Panel B), both measured at age 30. In columns (1) to (4) the conditioning set is gradually expanded. Individuals who were diagnosed with T1D in childhood earn DKK 34,422 less (column 4). This corresponds to a T1D income penalty of 14%. The likelihood of having positive labor income is 8.7 percentage points (pp) lower in the T1D group, which corresponds to a relative difference of 11%. It is worth noting that adding more control variables does not affect the estimated treatment effects much but increases precision slightly.

We next put this into context of the related literature that estimates the impact of a childhood shock. However, the related literature has largely focused on in-utero exposure and

[8] E.g., let $X_1 \in (X_A, X_B)$. Then the average difference, δ_{X_1} in treatment effect between the two groups is simply: $\delta_{X_1} = \frac{1}{|i:X_1=X_A|}\sum_{i:X_1=X_A} \hat{\tau}_i(X_1 = X_A) - \frac{1}{|i:X_1=X_B|}\sum_{i:X_1=X_B} \hat{\tau}_i(X_1 = X_B)$

very early life exposure while this paper focuses on a health shock that can occur throughout childhood. The most comparable setting would be that in Schwandt (2018), which finds that in-utero maternal exposure to influenza in Denmark leads to reductions in adult earnings of 9 percent, which is 64% of our estimated effect of 14 percent. Nilsson (2017) uses variation in Swedish alcohol policy to find that those most exposed to alcohol in utero and born to mothers under the age of 21 had 20 percent lower earnings which is 1.4 times the magnitude of our effect size.

Related literature has also found effects of childhood shocks outside of Scandinavia, although comparisons with the United States may be less than ideal. Despite the differences in settings, our results are similar in magnitude to Adhvaryu, Bednar, Molina, Nguyen, and Nyshadham (2020), which finds that the rollout of iodized salt across the United States led to an increase in income of 11 percent. Our results are also similar in magnitude to Smith (2009) which finds that controlling for education, which the author notes is negatively impacted by health, leads to income being 13 percent higher for those reporting their health to be excellent or very good. Our results are larger in magnitude to those in Beach, Ferrie, Saavedra, and Troesken (2016), which finds that eliminating early life exposure to typhoid fever led to a nine percent increase in earnings using an IV strategy. Our results are also larger than Isen, Rossin-Slater, and Walker (2017) which find the Clean Air Act of 1970 led to a one percent increase in annual earnings for the affected cohorts. Our estimate for labor force participation is about 13 times as large as the estimates in Adhvaryu et al. (2020) and 3 times as large as in Saez (2021), which studies the reduction in exposure to pneumonia in infancy in Chile.

To test for heterogeneous treatment effects, we next turn to our causal forest analysis. We graphically show the distribution of the estimated CATEs for both main outcomes in Figure 3. Panel A shows the distribution for labor income. The solid red line represents the mean of the conditional treatment effects whereas the dotted blue line is the OLS estimate (or

homogeneous treatment effect). The OLS estimate and the mean of the CATEs are not significantly different. The variation in the CATEs is considerable: the 10th percentile is DKK -47,949 and the 90th percentile is DKK -18,067. This is also true when looking at the likelihood of having positive labor income (-0.13 and -0.05 for the 10th and 90th percentile respectively; see Figure 3 Panel B). These results suggest that some subgroups of children are much less impacted by childhood onset T1D; thus, some parents can mitigate the effect.

To gain insight into the characteristics that are associated with the treatment heterogeneity, we proceed by grouping the sample by quartiles of the predicted treatment effects. We then compare the observable characteristics of those with larger predicted effects (Q1) vs. those with smaller predicted effects (Q4). These results are reported in Table 4. Panel A shows the mean characteristics from this analysis by labor market income. Being first born is associated with larger predicted penalties (the mean share of firstborns who are in Q1 is 0.50 vs. 0.41 in Q4). This larger impact on first born children is consistent with the larger educational spillovers stemming from childhood onset T1D for younger siblings than older siblings found in Eriksen et al. (2022). Among those with the highest predicted effects the ratio of males to females is roughly the same, however, males are much more likely to have smaller predicted income penalties.

Turning to the maternal characteristics, we observe that children of less educated mothers are more likely to have larger predicted effects. The share of mothers with only primary school as highest attained education is 0.64 in Q1 vs. 0.22 in Q4 – the share of mothers with medium tertiary education in Q1 is 0.03 vs. 0.45 in Q4. An age gradient is also at play: larger effects are associated with younger mothers. Larger predicted effects are also correlated with lower maternal income. In general, the paternal characteristics follow the same patterns as the maternal characteristics, although the educational gradient is less clear for the lowest levels of educational attainment.

The same pattern broadly emerges when looking at the CATEs regarding the outcome 'any labor market income': Males are underrepresented among those with numerically smaller penalties (share of males 0.43 vs. 0.66 in Q1 vs. Q4). We do note that here, the share of firstborns is large for those with smaller treatment effects, which seems inconsistent with the findings from the results on labor income. Other than that, the same SES gradient as for the other outcomes is observed. For example, the share of children where the mother only completed primary school is 0.61 among those with largest effects and 0.26 among those with smallest effects.

Of course, all these comparisons of characteristics are partial in that we are not holding all the other characteristics fixed. Having documented a SES gradient in the estimated heterogeneous treatment effects, it is worth noting that these treatment effects are in levels, and that the baseline mean labor market income or probability of being employed can vary considerably across subgroups. Thus, we proceed by quantifying the estimated CATEs relative to these baselines by selected maternal and paternal characteristics.

5.2 Relative Effects by Parental Socioeconomic Status

To further investigate whether high socioeconomic status parents can mitigate the impact we compare the estimated CATEs across maternal and paternal background characteristics, see Figures 4 and 5 (maternal and paternal). Looking at maternal education we see an inverse-U shaped pattern such that more education leads to smaller income penalties, except children of mothers with at least a master's degree (long tertiary education) have a larger penalty than children of mothers with a bachelor's degree. The mean CATE for labor income among individuals where the mother has no more than a primary education is DKK -37,984 while it is DKK -34,201 when the mother has a long tertiary education. While these means are statistically different, the difference in magnitude seems minor at the outset. However, taken relative to the income mean in each group this corresponds to relative effects of -17.8

percent and -13.5 percent respectively (Figure 4, Panel A on the left). We find a similar pattern with maternal educational background and having positive labor income, but a strong positive pattern for completed schooling beyond high school. The mean CATE for having any labor income is -10 percent relative to the mean when the mother has a long tertiary education and -12.7 percent when the mother has only primary school (Figure 4, Panel A).

Eriksen et al. (2021) examine the impact of a childhood T1D diagnosis on parental labor supply in Denmark. They find large impacts around the time of diagnosis for both parents, but mothers are more persistently impacted. Mothers experience a significant shift to part-time work through at least ten years after diagnosis. However, for mothers with at least a master's degree there is no impact on labor supply. This may explain the inverse-U shape we find for labor market outcomes.

When considering differences across the maternal income distribution we find the predicted earnings penalty is greatest for children of mothers who were in the bottom income quartile the year before the child's birth. At age 30, the mean CATE for these individuals is DKK -41,938 versus DKK -29,976 for those in the upper quartile. Relative to the mean this corresponds to -20% and -11% (Table 4, Panel A and Figure 4, Panel B to the left). This income gradient also holds for the probability of employment; see Figure 5.

We further investigate the relationship between maternal age at the time of the child's birth and the outcomes. The older the mother was at the child's birth, the less detrimental effects on the outcomes. The mean differences in CATEs across the age distribution are on par with those found for maternal income and education. Again, one should remember that when investigating these characteristics in isolation, we are not holding the other characteristics fixed. Consequently, having a mother or father of older age might also be correlated with parental educational attainment.

We similarly compare relative effects across paternal educational levels, paternal income quartile and paternal age at child's birth (Figure 4 and 5, right side). There is evidence of a strong gradient in paternal education and paternal income quartile for our main outcomes of interest. The key exception is that fathers who completed vocational training as their highest level of education have children with the second worse income penalties. The pattern for paternal age is more varied.

To summarize our findings thus far, we have clearly demonstrated a social gradient in the impact of childhood onset T1D. At age 30, individuals who were diagnosed with T1D in childhood have larger penalties if they were more disadvantaged to begin with, i.e., their mother or father was in a lower income quartile. Thus, we find that children are differentially impacted by family background, despite Currie and Stabile (2003) not finding evidence to support this as the mechanism behind the income health gradient in childhood. This difference may be due to Currie and Stabile (2003) using data from a different country (Canada) and health status being based on the health rating of the person most knowledgeable about the child (on a scale from 1-5), as opposed to focusing on a specific diagnosis.

5.3 Evidence of heterogeneity

Although the evidence on the associations between the estimated heterogeneity and the observable characteristics presented above is intuitive and meaningful, it is worthwhile to investigate further if the heterogeneity we observe is in fact true heterogeneity. We start out by estimating the group average treatment effect for those who are predicted to be in Q1 of the distribution. Among individuals predicted to be in Q1, the average treatment effect (S.E) is DKK -40,167.01 (9,174.26) vs. DKK -19,827.56 (9,589.63) in Q4 for the labor market income outcome. For the outcome 'any labor market income', we have -16.07 (2.11) pp. and -5.90 (1.77) pp. for Q1 and Q4, respectively. These are economically significant differences, but to formally test if the differences are significant in a statistical sense, we run regressions

with the two outcomes regressed on a set of dummy variables that indicates the quartile of predicted CATE (with Q1 as the left-out category). The results from this exercise are found in Table 5. For the labor market income, the treatment effects estimated within each group are not different in a statistical sense, as the CIs are overlapping. Again, it is worth mentioning that the difference between Q1 and Q4 is estimated at slightly more than DKK 20K. This corresponds to two thirds of the overall average treatment effect. For the outcome 'any labor market income', we observe statistically significant differences between Q1 and Q4. Individuals predicted to be in Q4 are on average 10 pp. more likely to have positive labor income compared with those in Q1. Again, this should be seen in relation to an overall average treatment effect of around 9 pp.

6 Mechanisms

6.1 Diabetes Management Outcomes

A potential mechanism behind the heterogeneity in adult labor market outcomes could be health capital. Previous literature has documented clinically meaningful differences in glucose control among children by socioeconomic status in Denmark. For example, Nielsen et al. (2019) document large gaps in glycated hemoglobin ($HbA_{1c)}$ levels by maternal education. Specifically, they find that while children have similar HbA_{1c} levels at the time of diagnosis, the gap becomes statistically significant starting two years after diagnosis. This matters because a large literature, starting with DCCT (1993), have shown a very strong relationship between higher HbA_{1c} (i.e., higher average glucose concentrations) and the risk of diabetes related complications such as retinopathy (leading to blindness), nephropathy (leading to kidney failure), and cardiovascular disease. The socioeconomic gradient in child health could translate to a socioeconomic gradient in adult health as well.

In Figure 6 we plot the mean group differences by maternal and paternal education, income quartile, and age at child's birth for HbA_{1c}. The left side of Panel A shows the mean group differences by maternal education. While the estimates are not precise, the general pattern shows that those with more highly educated mothers have better (lower) HbA_{1c} levels as adults[9]. Our results suggest that the gaps found in Nielsen et al. (2019) follow children into adulthood. Thus, children with less educated mothers may have larger long-run penalties because of worse disease management throughout childhood. The right side of Panel A shows the mean group differences by paternal education. The pattern is fairly similar to that of mothers.

Panel B shows a similar pattern by maternal (left) and paternal (right) income quartile. While the effects are not statistically different from each other, the results suggest better glucose control (and thus reduced risk of complications) for children with parents who are higher in the income distribution. Panel C shows a rather similar pattern by maternal (left) and paternal (right) age at child's birth.

In Figure 7 we plot the mean differences in specialty ambulatory care, again by maternal and paternal education, income quartile and age at child's birth. In Denmark, adults with T1D can receive care from their primary care physicians or at adult endocrinology clinics by diabetes specialists. Given that specialists should have more training and knowledge of T1D we may expect better outcomes for those who continue to receive medical care from a specialist as an adult. We can neither rule out differences across the groups, nor provide evidence of a clear pattern across the maternal or paternal characteristics.

In Figure 8 we plot mean differences in hospitalizations related to diabetic coma or diabetic ketoacidosis (DKA). While there are again no significant differences across the

[9] In practice, we use test values in a window of 5 years around the year the individuals turn 30. If more than one test is observed, we use the average. We observe at least one HbA_{1c} for 79% of the individuals with T1D.

groups, there is a clear pattern of more educated mothers and fathers having children who are less likely to be hospitalized in adulthood (Panel A). Panel B shows the impacts by maternal and paternal income quartile and the results also suggest better outcomes for more advantaged children. Lastly, we again find some suggestive evidence of better outcomes the older the mother and father are at child's birth (Panel C).

Lastly, in Figure 9 we plot mean differences in late diabetes complications by maternal and paternal education, income quartile and age at child's birth. While the standard errors overlap, we do find some evidence to suggest that having a more educated mother or father leads one to be less likely to develop any late complication by age 30. There are fewer clear trends for maternal and paternal income quartile and age at birth.

We more formally test for differences in health outcomes in adulthood by comparing the means in Q1 and Q4. Results are shown in the top portion of Table 6. The largest difference in characteristics for labor market income is for being hospitalized with diabetic coma or DKA. Forty three percent of those in Q1 have been hospitalized with diabetic coma or DKA, and 34 percent in Q4 have been, indicating a 9 pp. gap between the highest and lowest quartile. The two health outcomes with the largest difference across Q1 and Q4 for 'any labor market income' are diabetic coma or DKA, and HbA1c levels. Forty four percent of those in Q1 have been hospitalized with diabetic coma or DKA, and 36 percent in Q4 have, resulting in an 8 pp. gap. Individuals in Q1 have a mean HbA1c of 8.33 while those in Q4 have a mean HbA1c of 8.13, indicating worse glucose control among those in Q1 than Q4. We do not observe any difference across Q1 vs. Q4 on other health and utilization metrics, including LDL cholesterol levels, the probability of having ever smoked, and the probability of being treated for high cholesterol or high blood pressure.

Taken together these results provide suggestive evidence that both maternal and paternal socioeconomic status matters for T1D health related outcomes in adulthood and that better

adult health is tied to better adult labor market outcomes. Children of more educated mothers and children of mothers in higher income quartiles may have better disease management as proxied by HbA_{1c} levels and hospitalizations due to diabetic coma or DKA. We find a similar suggestive pattern for paternal education and paternal income quartile. Finding that parental characteristics matter for child health is consistent with medical literature documenting the role of parental behavior in children's glucose control (for example, see Davis, Delamater, Shawn, La Greca, Eidson, Perez-Rodiguez and Nembery (2001) and Thompson, Auslander and White (2001)).

6.2 Educational Outcomes

Since health capital impacts human capital (for example, see Bhalotra, and Venkataramani (2015), Karbownik and Wray (2019) and Saez (2021)) we next discuss how T1D impacts educational outcomes. Previous research has documented a clear relationship between glucose control and T1D educational penalties. For example, Skipper et al. (2019) show that children diagnosed with T1D have similar test scores compared with their peers on average, but those with worse glucose control have worse test scores and those with the best glucose control have better test scores. Eriksen et al. (2020) study the association of T1D and school wellbeing among middle school children and find similar outcomes (except children with T1D reported higher levels of headaches). However, those with worse glucose control again were found to have worse outcomes. Lindkvist et al. (2021) find that T1D is associated with worse 9^{th} grade exit exam scores, a higher relative risk of not completing 9^{th} grade by age 16 (end of compulsory school in Denmark), and a higher risk of not being enrolled in or graduated from upper secondary school by age 20. Again, the results indicate larger penalties for those with worse glucose control. Taken together, this research suggests that poorly controlled T1D in childhood can negatively impact a wide variety of educational outcomes.

Even those with well controlled T1D could be negatively impacted in the longer term through increased absenteeism. For example, Thingholm et al. (2020) find after diagnosis there is an increase in school absenteeism of around 50 percent more than the matched comparison children. Even children with well controlled glucose must miss school to attend medical appointments at the pediatric endocrinologist multiple times a year. There may be cumulative effects of consistently missing more school than peers without T1D do.

We next directly test for the role of educational attainment in explaining the labor market results. For each labor market outcome, we compare the mean of different levels of educational attainment for Q1 versus Q4. Results are shown in the bottom portion of Table 6. For labor market income, we find evidence of higher levels of educational attainment in Q4 than Q1. For example, 22 percent of those in Q1 have only attained primary school, while only 14 percent in Q4 have attained primary school. Eighteen percent of those in Q4 have attained a long tertiary education while only nine percent of those in Q1 have attained a long tertiary education. Similarly, for 'any labor market income' we find evidence of higher levels of educational attainment among those in Q4 than in Q1. For example, 23 percent of those in Q1 have only attained a primary education while only 14 percent of those in Q4 have attained a primary education. In terms of long tertiary education, only nine percent of those in Q1 have attained this level of education, while 17 percent of those in Q4 have attained it. Altogether, these results suggest impacts on educational attainment is likely an important mechanism in driving the differences in labor market outcomes.

6.3 Disease duration

As a last potential mechanism, we investigate if disease duration, i.e., age at onset of T1D, is associated with worse predicted outcomes. A fundamental challenge is that early onset is mechanically tied to being diagnosed early in the available data window we are using. As there have been large advancements in T1D management, such as new insulin

analogs, modern continuous glucose monitors and advancements in insulin pumps, we want to make sure that we are identifying the effects of duration and not picking up changes over time in treatment options. To address this, we create dummy variables for 'early onset' (age <9) and 'early time' (diagnosis year <1991) and regress them, along with their interaction, on the probability of being predicted to be in Q1 vs. Q4 (having numerically larger treatment effects) on the outcomes; see Table 7. For both our outcomes, the coefficient for early onset is positive and statistically significant, and the interaction with early time is also positive and significant. We interpret this as evidence that earlier onset of T1D is associated with worse outcomes.

7 Conclusion

In this paper we use the onset of T1D in childhood and Danish administrative registry data to study the impacts of a childhood health shock on adult labor market outcomes and test whether there is a socioeconomic gradient in the effects. We find significant and negative homogenous impacts on adult employment and labor income using OLS. However, this masks economically meaningful variation in the CATEs from the causal forest analysis. We find evidence of a socioeconomic gradient in the long-run impacts, in contrast with Currie and Stabile (2003) which concluded that the health-income gradient is not due to childhood health shocks differentially impacting children from high and low socioeconomic status families. Our results suggest that having a more highly educated mother or father or having a mother or father who is in a higher income quartile leads to smaller T1D penalties in adulthood.

It should be noted that conducting a similar analysis using data from another country may result in a different distribution of CATEs. Variation in access to health insurance, availability and access to endocrinology care, costs of medicine and medical devices, worker and student protections and anti-discrimination laws related to health and disability status

could all impact the results. However, our results do suggest that it is important that research using data from other countries explore the full distribution of impacts.

Future work could further explore the variation in parental responses in terms of investing in the child's health capital. Given the universal access to healthcare and low costs of treatment in Denmark, budget constraints seem like an unlikely explanation as to why differences in adult health outcomes exist. If it is due to differences in disease management knowledge or an inability to follow prescribed treatment plans, then maybe children with less educated mothers or fathers and children of lower income mothers and fathers would benefit from increased clinic visits, telehealth meetings or at-home support. Differences in parental preferences or discount rates may also be driving the effects, in which case rules surrounding when additional clinical interventions are implemented may need to be reevaluated.

References:

Athey, S., and Imbens, G. (2016), "Recursive Partitioning for Heterogeneous Causal Effects," Proceedings of the National Academy of Sciences, 113, 7353–7360.

Adhvaryu, Bednar, Molina, Nguyen, Nyshadham, (2020). "When It Rains It Pours: The Long-Run Economic Impacts of Salt Iodization in the United States," The Review of Economics and Statistics, vol 102(2), pages 395-407.

Almond, Currie, and Duque, (2018). "Childhood Circumstances and Adult Outcomes: Act II," Journal of Economic Literature, 56(4), 1360–1446

Beach, Brian, Joseph Ferrie, Martin Saavedra, and Werner Troesken, (2016). "Typhoid Fever, Water Quality, and Human Capital Formation." Journal of Economic History 76 (1): 41–75.

Becker, Gary and Tomes, Nigel, (1986), Human Capital and the Rise and Fall of Families, *Journal of Labor Economics*, 4, issue 3, p. S1-39.

Bhalotra, Sonia, and Atheendar Venkataramani. (2015). "Shadows of the Captain of the Men of Death: Health Innovation, Human Capital Investment, and Institutions." Working paper.

Breiman, L. (2001). Random forests. Machine learning, 45(1):5-32.

Bowles and Gintis (2002). The Inheritance of Inequality. *Journal of Economic Perspectives* Vol 16, Number 3, p3-30

Case, Lubotsky and Paxon (2002). Economic Status and Health in Childhood: The Origins of the Gradient. *The American Economic Review.* Vol. 92, No. 5, p1334-1308

Carneiro, Pedro, Katrine V. Løken, and Kjell G. Sal-vanes. (2015). "A Flying Start? Maternity Leave Ben-efits and Long-Run Outcomes of Children." Journal of Political Economy 123 (2): 365–412.

Cobry, E and Jaser, S. (2019). Brief Literature Review: The Potential of Diabetes Technology to Improve Sleep in Youth With Type 1 Diabetes and Their Parents: An Unanticipated Benefit of Hybrid Closed-Loop Insulin Delivery Systems. *Diabetes Spectrum* DOI**:** 10.2337/ds18-0098

Currie and Stabile (2003). Socioeconomic Status and Child Health: Why Is the Relationship Stronger for Older Children? *The American Economics Review.* Vol 93, No. 5, p 1813-1823

Davis, Delamater, Shaw, La Greca, Eidson, Perez-Rogdriguez, Nemery. (2001). Parenting styles, regimen adherence, and glycemic control in 4- to 10-year-old children with diabetes. *J Pediatric Psychology* Mar;26(2):123-9. doi: 10.1093/jpepsy/26.2.123

Datar A, Rebecca KM, Loughran DS (2010) Endowments and parental investments in infancy and early childhood. Demography 47(1):145–162

DCCT 1993 (The Diabetes Control and Complications Trial Research Group), (1993) The Effect of Intensive Treatment of Diabetes of the Development and Progression of Long-Term Complications in Insulin-Dependent Diabetes Mellitus, *New England Journal of Medicine*, vol. 329: 977-986.

Eriksen, Gaulke, Skipper and Svensson (2021). The Impact of Childhood Health Shocks on Parental Labor Supply. *Journal of Health Economics*. Vol 78

Eriksen, Gaulke, Thingholm, Svensson and Skipper (2020). Association of type 1 diabetes and school wellbeing: a population-based cohort study of 436,439 Danish schoolchildren. *Diabetologia*. https://doi.org/10.1007/s00125-020-05251-z

Eriksen, Gaulke, Thingholm, Svensson and Skipper (2022). Educational Consequences of Sibling's Disability. Working paper.

Grätz, Michael and Torche, Florencia. (2016). Compensation or reinforcement? The stratification of parental responses to children's early ability. Demography, 53(6):1883–1904.

Gensowski, Nielsen, Nete M. Nielsen, Rossin-Slater and Wüst, (2019). "Childhood health shocks, comparative advantage, and long-term outcomes: Evidence from the last Danish polio epidemic". Journal of Health Economics, volume 66, pages 27-36 https://doi.org/10.1016/j.jhealeco.2019.03.010

Guo and Zhang, (2021). Altruistic or Exchange Motive? Evidence on the effect of children's health shocks on intra-household resource allocation. Working paper.

Hastie, Trevor, Robert Tibshirani and Jerome Friedman, (2009). *Elements of Statistical Learning: Data Mining, Inference and Prediction (Second Edition).*

Havnes, Tarjei, and Magne Mogstad. (2011). "No Child Left Behind: Subsidized Child Care and Children's Long-Run Outcomes." American Economic Journal: Economic Policy 3 (2): 97–129.

Havnes, Tarjei, and Magne Mogstad. (2015). "Is Universal Child Care Leveling the Playing Field?" Journal of Public Economics 127: 100–114.

Houmark, Ronda and Rosholm, (2020). "The Nurture of Nature and the Nature of Nurture: How Genes and Investments Interact in the Formation of Skills". IZA Discussion Papers No. 13780.

Hsin A. (2012). Is biology destiny? Birth weight and differential parental treatment. *Demography*, *49*(4), 1385–1405. https://doi.org/10.1007/s13524-012-0123-y

Isen, Adam, Maya Rossin-Slater, and W. Reed Walker. (2017). "Every Breath You Take—Every Dollar You'll Make: The Long-Term Consequences of the Clean Air Act of 1970." Journal of Political Economy 125(3): 848–902.

Karbownik, Krzysztof and Anthony Wray, (2019). "Educational, Labor-market and Intergenerational Consequences of Poor Childhood Health" NBER Working Paper No. 26368

Lindkvist, Thorsen, Paulsrud, Thingholm, Eriksen, Gaulke, Skipper and Svensson (2021). Association of type 1 diabetes and educational achievement in 16–20-year-olds: A Danish nationwide register study. *Diabetic Medicine.* DOI: 10.1111/dme.14673

Løken, Katrine V., Magne Mogstad, and Matthew Wiswall. (2012). "What Linear Estimators Miss: The Effects of Family Income on Child Outcomes." American Economic Journal: Applied Economics 4(2):1–35.

Nielsen, Gaulke, Eriksen, Svensson, and Skipper (2019). Socioeconomic inequality in metabolic control among children with type 1 diabetes: a nationwide longitudinal study of 4,079 Danish children", *Diabetes Care*, 2019 Aug; 42(8): 1398-1405.

Nilsson, J. Peter, (2017). "Alcohol Availability, Prenatal Conditions, and Long-Term Economic Outcomes." Journal of Political Economy, 125 (4): 1149–207.

Persson, Gerdtham and Carlsson (2016). "Labor market consequences of childhood onset type 1 diabetes" Economics and Human Biology. 23, p. 180-192.

Pillar, Schuscheim, Weiss, Malhotra, McCowen, Shlitner, Peled, and Shehadeh. (2003). Interactions between hypoglycemia and sleep architecture in children with type 1 diabetes mellitus, *The Journal of Pediatrics*, Volume 142, Issue 2, Pages 163-168, https://doi.org/10.1067/mpd.2003.66.

Pociot and Lernmark. (2016). Genetic risk factors for type 1 diabetes, *The Lancet*, Volume 387, Issue 10035, Pages 2331-2339, https://doi.org/10.1016/S0140-6736(16)30582-7.

Regnell, S., and Lernmark, A. (2017). Early Prediction of Autoimmune (type 1) diabetes. *Diabetologia*, 1370-1381.

Restrepo, Brandon J. (2016). Parental investment responses to a low birth weight outcome: Who compensates and who reinforces? Journal of Population Economics, 29(4):969–989.

Saez, Camila. (2021). Long-term effects of early life exposure to pneumonia. Working paper.

Schwandt, Hannes, (2008). "The Lasting Legacy of Seasonal Influenza: In-Utero Exposure and Labor Market Outcomes." Centre for Economic Policy Research Discussion Paper DP12563

Smith (2009). The Impact of Childhood Health on Adult Labor Market Outcomes. *The Review of Economics and Statistics*, August, 91(3): 478–489

Skipper, N., Gaulke, A., Sildorf, S.M., Eriksen, T.M., Nielsen, N.F., Svensson, J. (2019). Association of Type 1 Diabetes with Standardized Test Scores of Danish Schoolchildren. *JAMA.* 321(5), 484–492. doi:10.1001/jama.2018.21819

Tao, B., Pietropaolo, M., Atkinson, M., Schatz, D., and Taylor, D. (2010). Estimating the cost of Type 1 Diabetes in the US: A Propensity Score Matching Method. *Plos ONE*, July 2010

Thingholm, P. R., Gaulke, A., Eriksen, T. M., Svensson, J., Skipper, N. (2020). Association of Prodromal Type 1 Diabetes With School Absenteeism of Danish Schoolchildren: A

Population-Based Case-Control Study of 1,338 Newly Diagnosed Children, *Diabetes Care*, Nov. 43(11), 2886-2888, 2020,

Thompson, Auslander and White (2001). Comparison of Single-Mother and Two-Parent Families on Metabolic Control of Children with Diabetes. *Diabetes Care*;24(2):234–238 https://doi.org/10.2337/diacare.24.2.234

Wager and Athey, (2018). Estimation and Inference of Heterogeneous Treatment Effects using Random Forests. Journal of the American Statistical Association, doi:10.1080/01621459.2017.1319839

Table 1: Descriptive statistics

		Diabetes	No diabetes
Number of observations		1810	9050
Child characteristics			
Birth order	First born (0/1)	0.45	0.47
	Second born (0/1)	0.40	0.38
	Third born or later (0/1)	0.16	0.15
Sex	Male (0/1)	0.53	0.51
Paternal Characteristics			
Education	Primary school (0/1)	0.33	0.31
	Secondary school (0/1)	0.03	0.03
	Vocational (0/1)	0.44	0.43
	Short Ter. (0/1)	0.02	0.03
	Medium Ter. (0/1)	0.11	0.13
	Long Ter. (0/1)	0.07	0.07
Income	First quartile (0/1)	0.17	0.17
	Second quartile (0/1)	0.16	0.17
	Third quartile (0/1)	0.16	0.17
	Fourth quartile (0/1)	0.17	0.16
Immigration status	Native (0/1)	0.97	0.97
	Immigrant or descendent (0/1)	0.03	0.03
Age	<25 (0/1)	0.07	0.09
	25-29 (0/1)	0.30	0.31
	30-34 (0/1)	0.35	0.34
	35+ (0/1)	0.27	0.26
Maternal Characteristics			
Education	Primary school (0/1)	0.41	0.39
	Secondary school (0/1)	0.04	0.03
	Vocational (0/1)	0.31	0.31
	Short Ter. (0/1)	0.03	0.03
	Medium Ter. (0/1)	0.18	0.20
	Long Ter. (0/1)	0.03	0.03
Income	First quartile (0/1)	0.18	0.17
	Second quartile (0/1)	0.16	0.17
	Third quartile (0/1)	0.17	0.17
	Fourth quartile (0/1)	0.16	0.17
Immigration status	Native (0/1)	0.98	0.98
	Immigrant or descendent (0/1)	0.02	0.02
Age	<25 (0/1)	0.21	0.22
	25-29 (0/1)	0.39	0.40
	30-34 (0/1)	0.28	0.27
	35+ (0/1)	0.12	0.11

Notes: Descriptive statistics of individuals diagnosed with diabetes before the age of 18 and a group matched on birthday and onset year.

Table 2: Linear regression of diabetes status on observable characteristics

		Coeff.	S.E.
Child characteristics			
Birth order	First born (0/1)	Ref.	
	Second born (0/1)	0.002	(0.009)
	Third born or later (0/1)	-0.009	(0.012)
Sex	Male (0/1)	0.014**	(0.007)
Paternal characteristics			
Education	Primary school (0/1)	Ref.	
	Secondary school (0/1)	0.002	(0.022)
	Vocational (0/1)	-0.004	(0.009)
	Short Ter. (0/1)	-0.035	(0.022)
	Medium Ter. (0/1)	-0.025*	(0.013)
	Long Ter. (0/1)	-0.016	(0.018)
Income	First quartile (0/1	Ref.	
	Second quartile (0/1)	-0.005	-0.013
	Third quartile (0/1)	-0.007	-0.013
	Fourth quartile (0/1)	0.008	-0.014
Immigration status	Immigrant or descendent (0/1)	-0.018	(0.022)
Age			
	<25 (0/1)	Ref.	
	25-29 (0/1)	0.028**	(0.014)
	30-34 (0/1)	0.035**	(0.015)
	35+ (0/1)	0.035**	(0.017)
Maternal characteristics			
Education	Primary school (0/1)	Ref.	
	Secondary school (0/1)	0.005	(0.022)
	Vocational (0/1)	-0.012	(0.009)
	Short Ter. (0/1)	-0.016	(0.021)
	Medium Ter. (0/1)	-0.026**	(0.011)
	Long Ter. (0/1)	-0.013	(0.023)
Income	First quartile (0/1)	Ref.	
	Second quartile (0/1)	-0.015	(0.013)
	Third quartile (0/1)	-0.015	(0.013)
	Fourth quartile (0/1)	-0.014	(0.014)
Immigration status	Immigrant or descendent (0/1)	-0.022	(0.024)
Age	<25 (0/1)	Ref.	
	25-29 (0/1)	0.006	-0.011
	30-34 (0/1)	0.017	-0.014
	35+ (0/1)	0.02	-0.017
	Observations	10,860	
	R-squared	0.003	
	Joint F-test	0.589	
	Prob > F	0.995	

Notes: OLS regression of diabetes status on observable characteristics. The OLS additionally controls for cohort, onset-year and municipality fixed effects. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Standard errors are clustered at the match group level.

Table 3: Average Treatment Effects of Diabetes onset in childhood.

Panel A	(1) Income (DKK)	(2) Income (DKK)	(3) Income (DKK)	(4) Income (DKK)
Diabetes	-33,102***	-32,978***	-32,453***	-32,422***
	(4,540)	(4,482)	(4,487)	(4,460)
Observations	10,418	10,418	10,418	10,418
Mean Outcome	239,468.45	239,468.45	239,468.45	239,468.45
Paternal Char.	No	Yes	No	Yes
Maternal Char.	No	No	Yes	Yes
R-squared	0.047	0.067	0.069	0.078

Panel B	(1) Any LMI (0/1)	(2) Any LMI (0/1)	(3) Any LMI (0/1)	(4) Any LMI (0/1)
Diabetes	-0.0875***	-0.0875***	-0.0866***	-0.0867***
	(0.0110)	(0.0109)	(0.0109)	(0.0109)
Observations	10,418	10,418	10,418	10,418
Mean Outcome	0.81	0.81	0.81	0.81
Paternal Char.	No	Yes	No	Yes
Maternal Char.	No	No	Yes	Yes
R-squared	0.043	0.054	0.058	0.064

Notes: Standard errors in parentheses. *** p<0.01. ** p<0.05. * p<0.1. Estimates of treatment impact of diabetes onset in childhood. The estimates are coefficients from separate OLS regressions of income, or an indicator for having any income. Column 1 controls for child characteristics, cohort, onset-year and municipality fixed effects. Column 2 and 3 additionally controls for paternal or maternal characteristics. Column 4 controls for both paternal and maternal characteristics. Mean outcomes are reported for the control group.

Table 4: Mean child, maternal, and paternal characteristics by predicted CATE

		Panel A: Labor market income				Panel B: Any labor market income			
Child characteristics		Q1	Q4	Difference	t stat	Q1	Q4	Difference	t stat
Birth order	First born (0/1)	0.50	0.41	0.10	7.00	0.34	0.57	-0.23	-16.89
	Second born (0/1)	0.37	0.46	-0.09	-6.73	0.54	0.29	0.25	19.00
	Third born or later (0/1)	0.13	0.14	-0.01	-0.53	0.12	0.15	-0.02	-2.35
Sex	Male (0/1)	0.49	0.58	-0.08	-6.15	0.43	0.66	-0.22	-16.72
Maternal characteristics									
Education	Primary school (0/1)	0.64	0.22	0.43	34.42	0.61	0.26	0.34	26.75
	Secondary school (0/1)	0.05	0.02	0.03	5.88	0.03	0.03	0.00	0.78
	Vocational (0/1)	0.22	0.27	-0.05	-3.78	0.23	0.35	-0.12	-9.46
	Short Ter. (0/1)	0.03	0.03	0.00	1.00	0.02	0.03	-0.01	-3.48
	Medium Ter. (0/1)	0.03	0.45	-0.42	-40.77	0.09	0.29	-0.20	-18.84
	Long Ter. (0/1)	0.03	0.02	0.01	1.27	0.02	0.04	-0.02	-3.75
Age	<25 (0/1)	0.28	0.13	0.14	12.96	0.26	0.22	0.04	3.58
	25-29 (0/1)	0.53	0.30	0.24	17.77	0.52	0.24	0.28	21.92
	30-34 (0/1)	0.14	0.43	-0.29	-24.20	0.16	0.39	-0.22	-18.73
	35+ (0/1)	0.05	0.14	-0.09	-11.45	0.05	0.15	-0.10	-12.08
Income	First quartile (0/1)	0.37	0.05	0.32	31.27	0.44	0.02	0.42	41.81
	Second quartile (0/1)	0.25	0.08	0.17	17.11	0.21	0.09	0.12	12.45
	Third quartile (0/1)	0.18	0.12	0.06	6.10	0.12	0.13	-0.01	-1.08
	Fourth quartile (0/1)	0.09	0.23	-0.14	-13.93	0.06	0.25	-0.19	-19.30
Immigration status	Immigrant or descendent (0/1)	0.02	0.03	-0.01	-2.02	0.02	0.03	0.00	-0.91
Paternal characteristics									
Education	Primary school (0/1)	0.33	0.35	-0.02	-1.51	0.41	0.24	0.17	13.21
	Secondary school (0/1)	0.03	0.04	-0.02	-3.23	0.03	0.03	0.00	0.64
	Vocational (0/1)	0.52	0.30	0.21	16.09	0.40	0.47	-0.06	-4.69
	Short Ter. (0/1)	0.03	0.02	0.00	1.06	0.03	0.03	0.00	0.59
	Medium Ter. (0/1)	0.07	0.18	-0.11	-12.32	0.08	0.14	-0.06	-6.53
	Long Ter. (0/1)	0.04	0.11	-0.07	-10.23	0.04	0.09	-0.05	-7.77
Age	<25 (0/1)	0.13	0.04	0.09	11.16	0.12	0.08	0.04	5.24
	25-29 (0/1)	0.39	0.19	0.19	15.88	0.43	0.21	0.22	17.84
	30-34 (0/1)	0.31	0.44	-0.13	-9.45	0.30	0.39	-0.10	-7.34
	35+ (0/1)	0.17	0.33	-0.16	-13.19	0.15	0.32	-0.17	-14.80
Income	First quartile (0/1	0.33	0.06	0.28	26.79	0.39	0.02	0.37	37.76
	Second quartile (0/1)	0.22	0.09	0.13	12.75	0.20	0.13	0.07	6.52
	Third quartile (0/1)	0.22	0.13	0.09	9.09	0.16	0.15	0.00	0.07
	Fourth quartile (0/1)	0.12	0.17	-0.05	-5.26	0.09	0.17	-0.08	-8.29
Immigration status	Immigrant or descendent (0/1)	0.02	0.03	-0.01	-3.05	0.03	0.03	0.00	0.34

quartile, Q1 vs. Q4

Notes: The table reports the means and differences in means for child, maternal, and paternal characteristics by predicted quartile of the conditional average treatment effects. Q1 represents the numerically largest effects and Q4 the smallest. Panel A shows the differences by labor market income and Panel B for the outcome 'any labor market income'.

Table 5: Group Average Treatment Effects (GATES) by predicted CATE quartile

	Labour Market Income				Any Labor Market Income		
	Difference	95% CI			Difference	95% CI	
Q1	Ref.			Q1	Ref.		
Q2	36.07	-25898.90	25971.04	Q2	0.10	0.05	0.15
Q3	4153.41	-21241.75	29548.57	Q3	0.09	0.03	0.14
Q4	20339.45	-5570.71	46249.62	Q4	0.10	0.05	0.15

Notes: The sample is partitioned into quartiles by the predicted conditional average treatment effect. Group average treatment effects are then estimated using a set of quartile dummy indicators. This is done for both outcomes.

Table 6: Differences in educational attainment and selected health outcomes by predicted treatment quartile

		Panel A: Labor market income				Panel B: Any labor market income			
		Q1	Q4	Difference	t stat	Q1	Q4	Difference	t stat
Diabetes Sample									
	Diabetic coma/DKA (0/1)	0.43	0.34	0.09	2.66	0.44	0.36	0.08	2.39
	Late complications (0/1)	0.55	0.53	0.02	0.57	0.57	0.57	0.00	0.03
	HbA1c	8.19	8.13	0.06	0.60	8.33	8.13	0.20	2.03
	LDL cholesterol	2.54	2.52	0.02	0.38	2.56	2.60	-0.04	-0.77
	Ever smoker (0/1)	0.32	0.34	-0.02	-0.67	0.36	0.31	0.05	1.69
	Hypertension treatment (0/1)	0.17	0.19	-0.02	-0.62	0.20	0.19	0.01	0.30
	Lipid lowering treatment (0/1)	0.14	0.14	0.01	0.38	0.17	0.14	0.03	1.23
Full Sample									
Education	Primary school (0/1)	0.22	0.14	0.08	7.82	0.23	0.14	0.09	8.75
	Secondary school (0/1)	0.07	0.09	-0.01	-1.80	0.08	0.08	0.00	-0.47
	Vocational (0/1)	0.38	0.29	0.09	6.98	0.37	0.33	0.04	3.33
	Short Ter. (0/1)	0.04	0.06	-0.01	-2.08	0.04	0.06	-0.02	-3.12
	Medium Ter. (0/1)	0.18	0.24	-0.06	-5.57	0.18	0.22	-0.04	-3.61
	Long Ter. (0/1)	0.09	0.18	-0.09	-9.47	0.09	0.17	-0.08	-8.41

Notes: Differences in the probability of selected health outcomes for individuals with T1D by treatment effect quartile. Differences in educational attainment for the full sample by treatment effect quartile. DKA is diabetic ketoacidosis.

Table 7: Probability of being in Q1 vs. Q4 by age at T1D onset

	Labor market income		Any labor market income	
	Coef	S.E.	Coef	S.E.
Early onset	0.18	0.03	0.15	0.03
Early onset x Early time	0.29	0.04	0.16	0.04

Notes: The probability of being in treatment effect quartile 1 vs. 4 by age at onset (< 9 years) and the interaction with early time (diagnosed before 1991).

Figure 1: Mean difference (95% CI) in days absent from school during a given month relative to diagnosis of type 1 diabetes (diabetes vs. no diabetes)

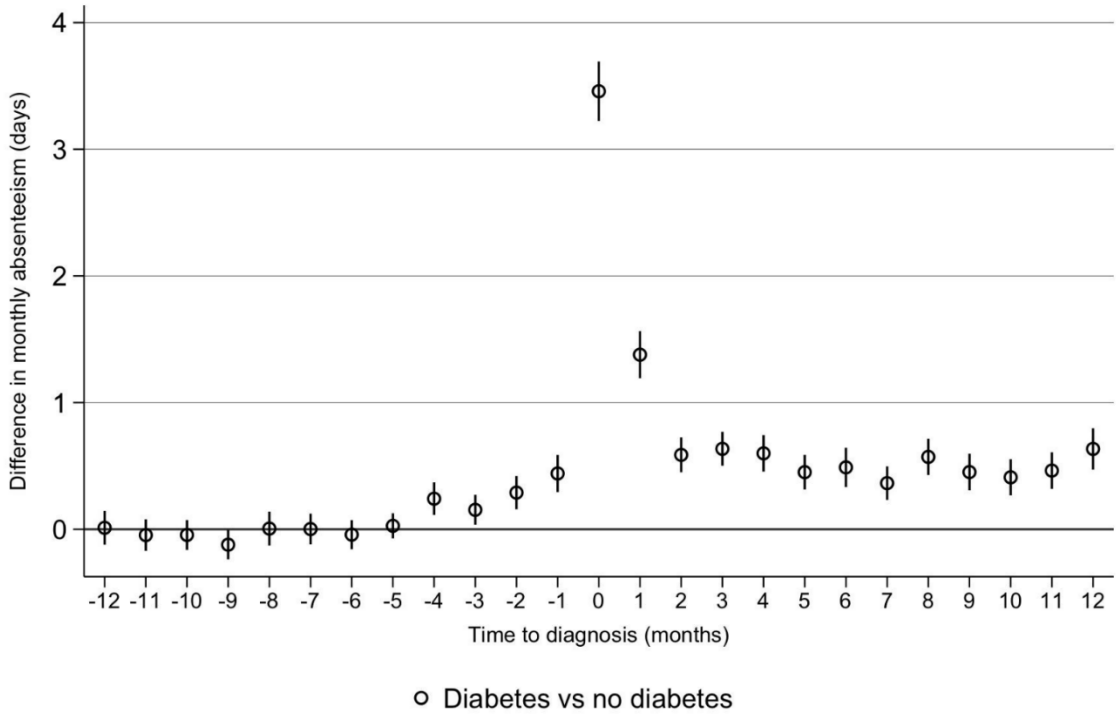

Notes: n= 1,338 children diagnosed with type 1 diabetes from August 1 2010 to June 30 2017 compared with n= 6,690 age and sex matched controls. Mean (95% CI) difference in number of days absent from school relative to diabetes diagnosis (month 0). The mean differences are adjusted for calendar-month and school grade specific effects. As the month of July is the only month of year with no school days in Denmark, it was left out of the analysis. Months -12 to -5 showed non-significant differences (with a level of significance at $p<0.05$).

Published previously in Thingholm et al. Association of Prodromal Type 1 Diabetes with School Absenteeism of Danish Schoolchildren: A Population-Based Case-Control Study of 1,338 Newly Diagnosed Children. Diabetes Care 2020 Nov; 43(11): 2886-2888. Copyright 2020 by the American Diabetes Association.

Figure 2: Differences in health care utilization around the time of diagnosis, treatment vs. control

Panel A: Probability of any hospital admission

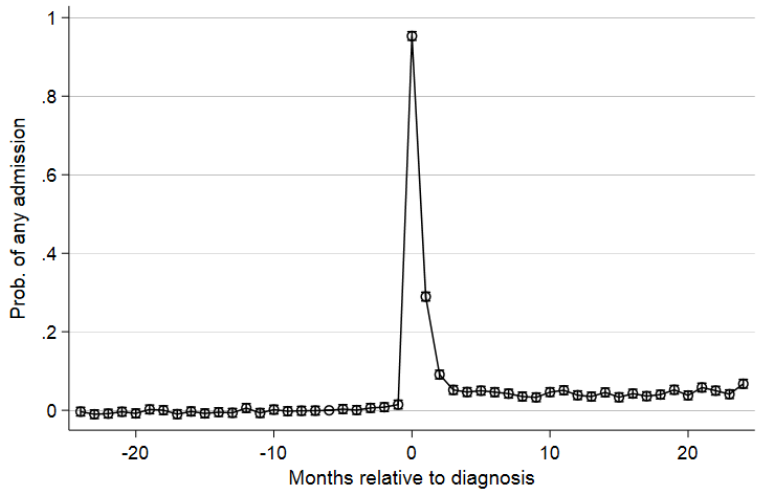

Panel B: Probability of any GP visit

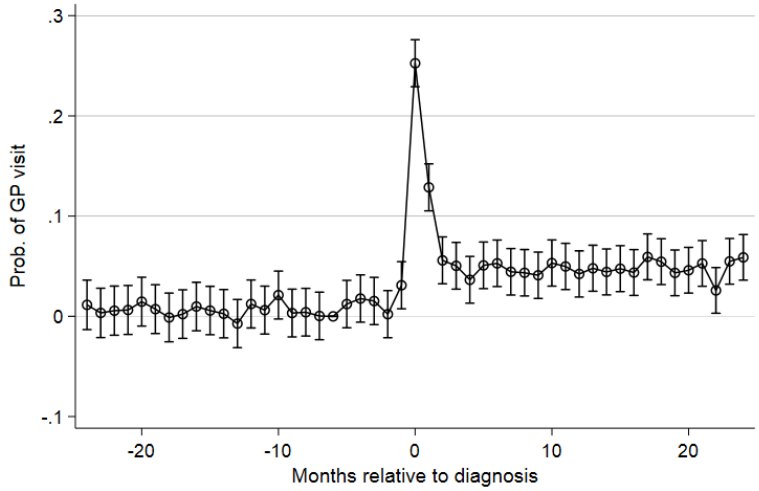

Panel C: Probability of any pharmacy claim

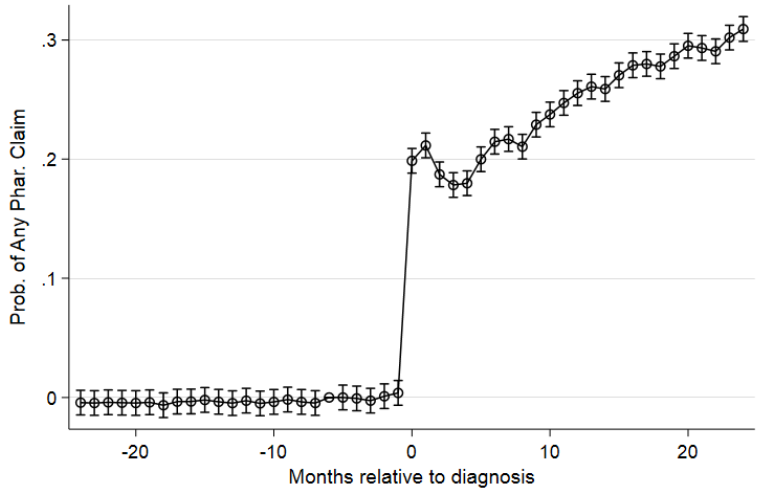

Notes: Panel A shows the difference in 'any hospital admission' by month relative to the diagnosis month for treatment vs comparison individuals. In Panel B, the outcome is the probability of visiting the general practitioner, and Panel C shows the difference in the probability of having a pharmacy claim.

Figure 3: Estimated distributions of Conditional Average Treatment Effects

Panel A: Estimated distribution of CATEs for income at age 30

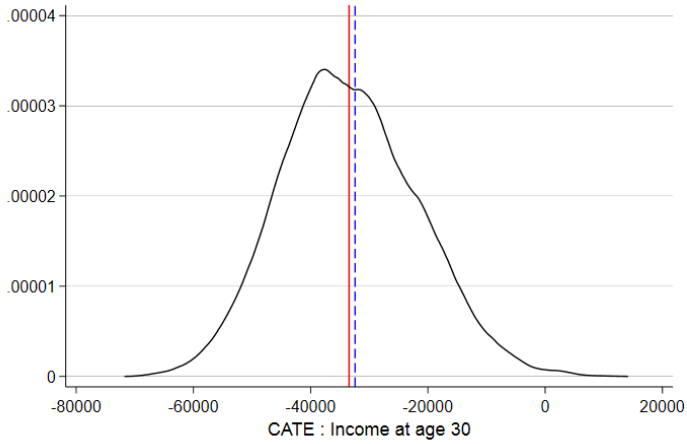

Panel B: Estimated distribution of CATEs for having positive LMI at age 30

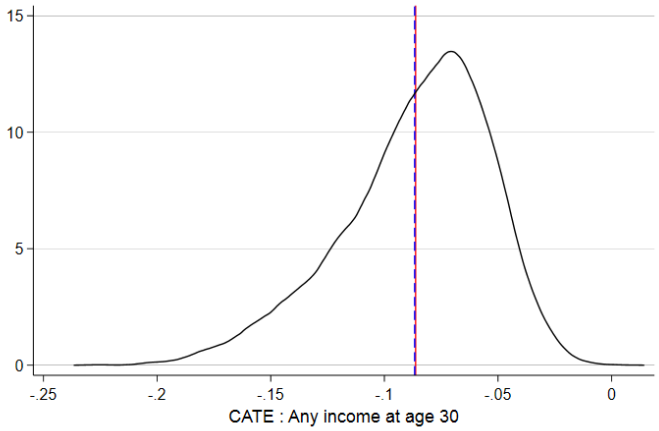

Notes: The solid (red) line indicates the average CATE in the sample. The dashed (blue) line is the treatment effect estimate from an OLS regression including the full set of controls (equal to the coefficient reported in column (4) in table 3). The average CATE is not statistically significantly different from the OLS treatment estimate for any of the outcomes. The estimated distributions are capped at the 1st and 99th percentiles.

Figure 4: Relative treatment effects on income at age 30 by parental characteristics

Panel A: Relative treatment effects by parental education

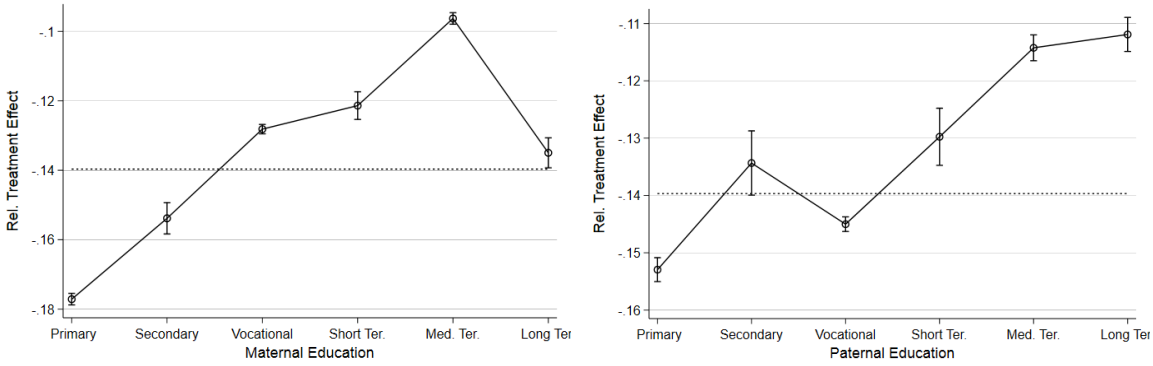

Panel B: Relative treatment effects by parental income quartile

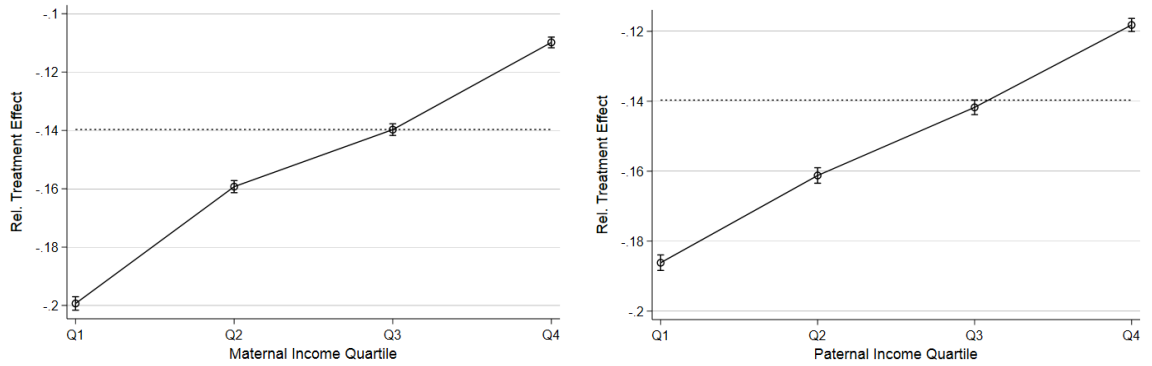

Panel C: Relative treatment effects by parental age

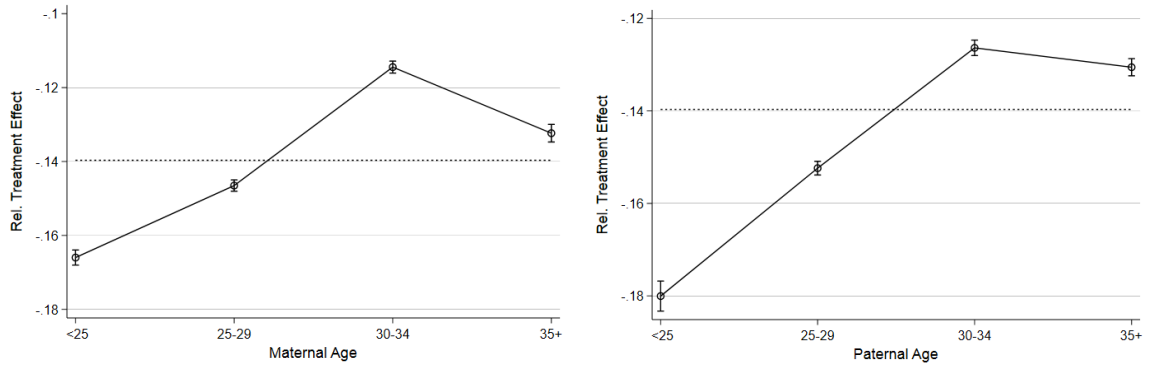

Notes: This figure presents subgroup effects and 95% CI for the impact of diabetes on income at age 30. The graph shows mean CATE from causal forest analysis by maternal education, age, and income quartile. The mean is scaled by the mean outcome in the subgroup. The horizontal line indicates the relative treatment effect for the entire population

Figure 5: Relative treatment effects on any labor market income at age 30 by maternal characteristics

Panel A: Relative treatment effects by parental education

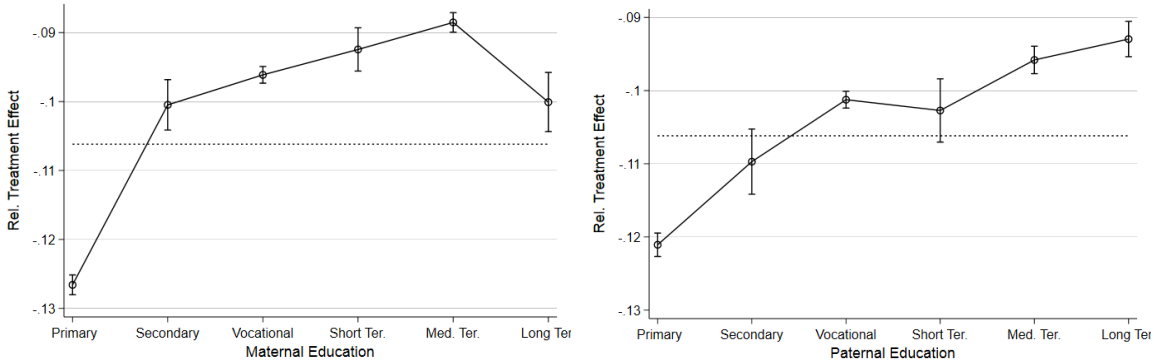

Panel B: Relative treatment effects by parental income quartile

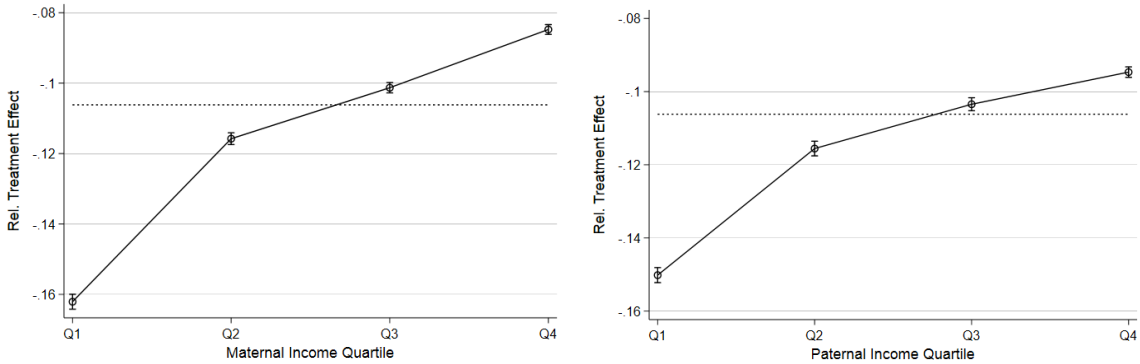

Panel C: Relative treatment effects by parental age

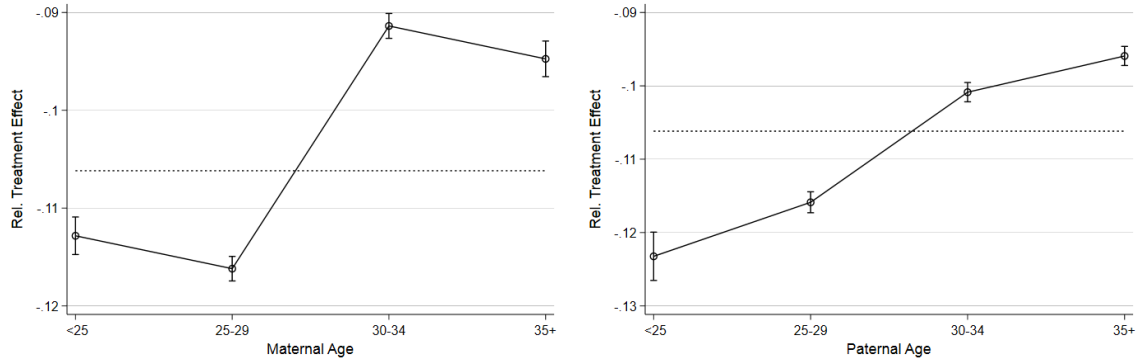

Notes: This figure presents subgroup effects and 95% CI for the impact of diabetes on the probability of having any labor market income at age 30. The graph shows mean CATE from causal forest analysis by maternal education, age, and income quartile. The mean is scaled by the mean outcome in the subgroup. The horizontal line indicates the relative treatment effect for the entire population

Figure 6: Difference in HbA_{1C} for People with T1D by parental characteristics

Panel A: Difference by parental education

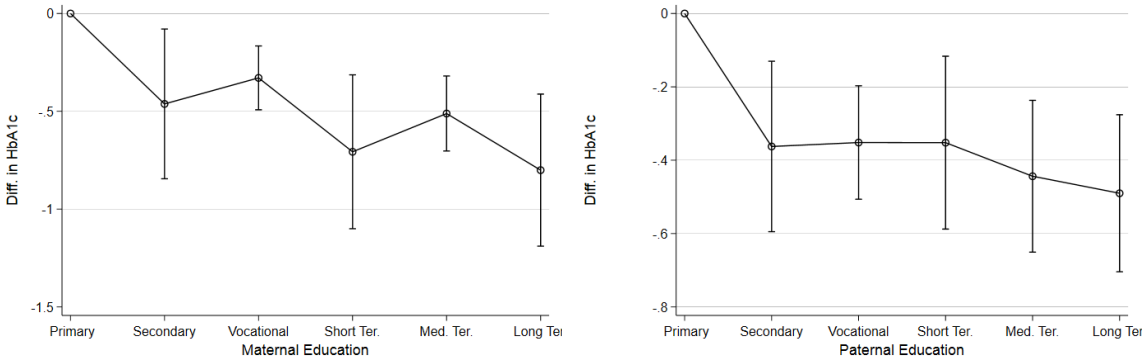

Panel B: Difference by parental income quartile

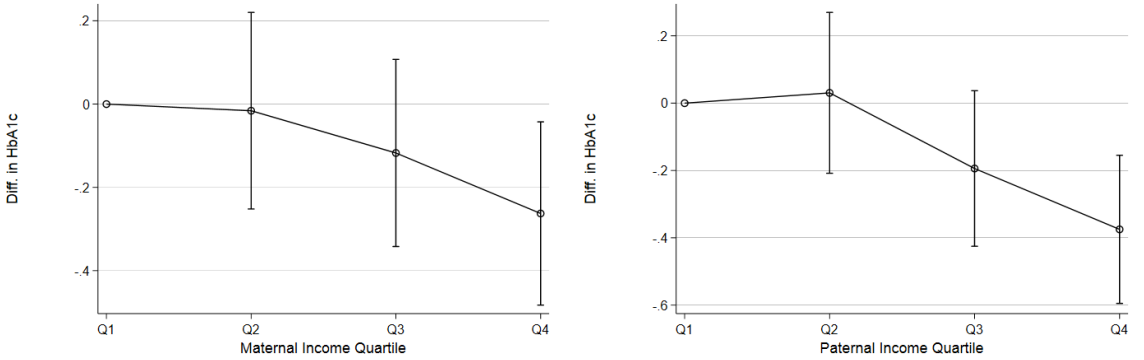

Panel C: Difference by parental age

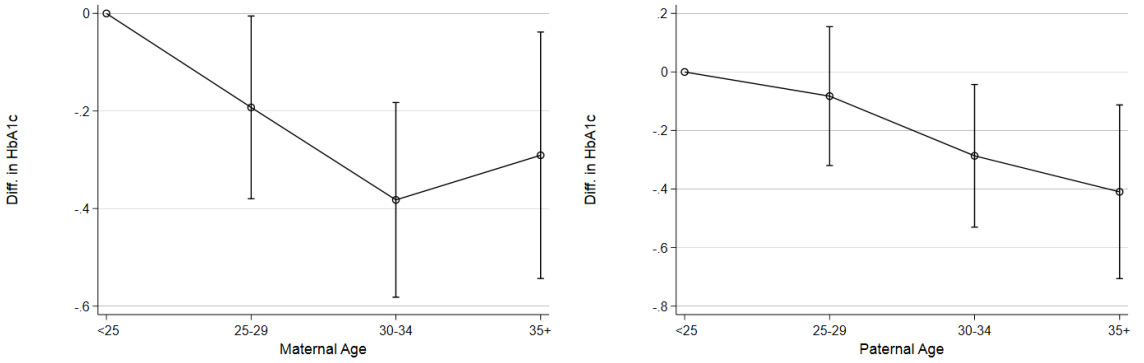

Notes: This figure presents subgroup differences in diabetes related outcomes among the individuals with diabetes. Glycated hemoglobin (HbA_{1c}) is a measure of how well the glucose levels are managed with lower values indicating better disease management. Mean differences relative to the comparison group are reported with 95% CI. The outcome mean is 8.2.

Figure 7: Difference in probability of receiving ambulatory care related to T1D by parental characteristics

Panel A: Difference by parental education

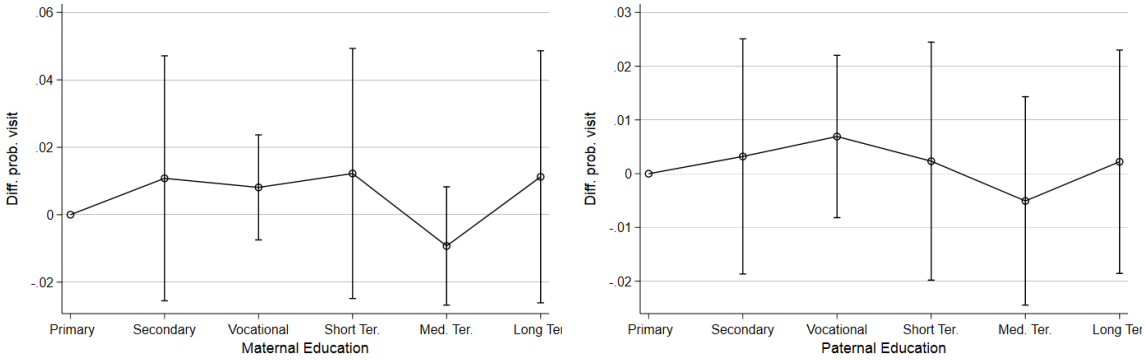

Panel B: Difference by parental income quartile

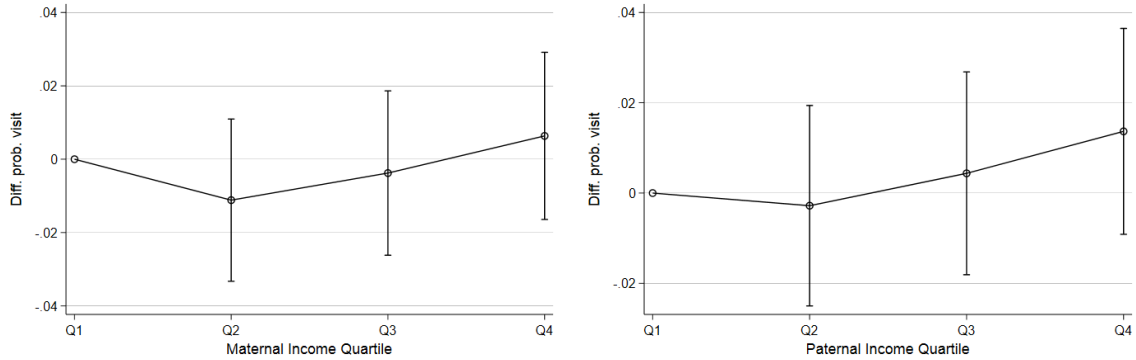

Panel C: Difference by parental age

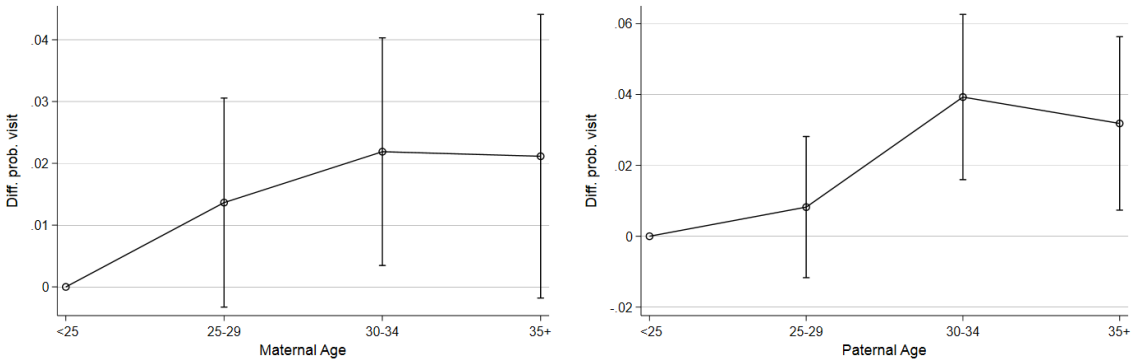

Notes: This figure presents subgroup differences in diabetes related outcomes among the individuals with diabetes. The outcome is the probability of receiving specialized ambulatory care. Mean differences relative to the comparison group are reported with 95% CI. The outcome mean is 0.77

Figure 8: Difference in probability of hospital admission with diabetes related acute conditions by parental characteristics

Panel A: Difference by parental education

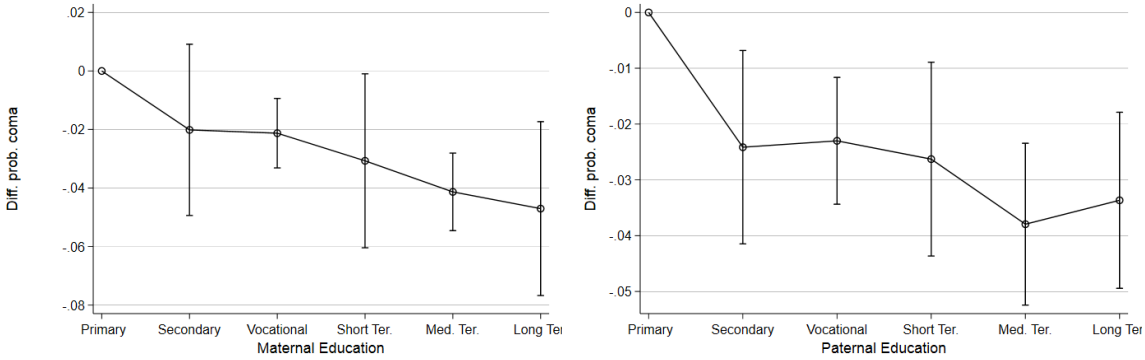

Panel B: Difference by parental income quartile

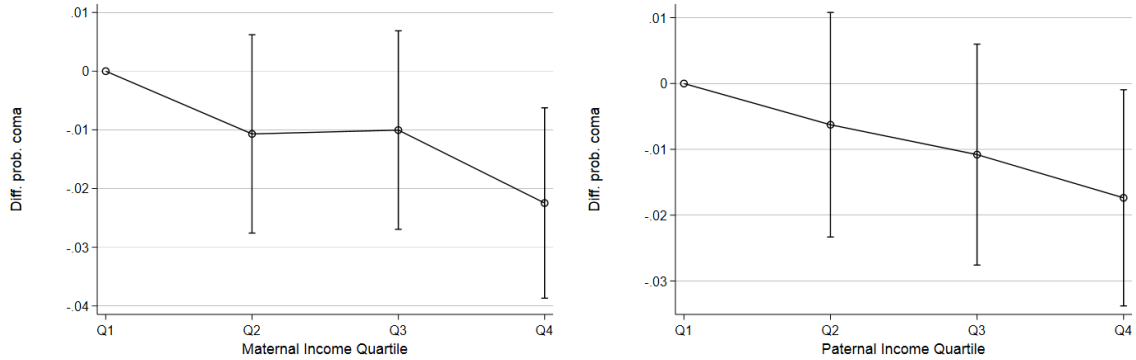

Panel C: Difference by parental age

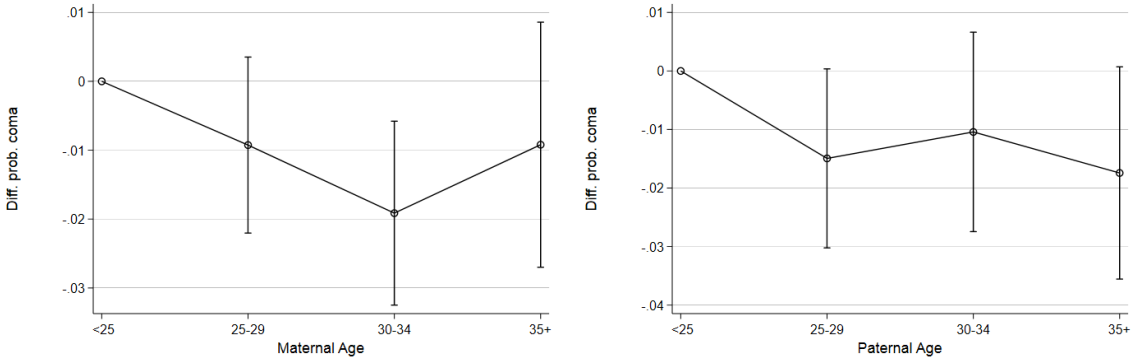

Notes: This figure presents subgroup differences in diabetes related outcomes among the individuals with diabetes. The outcome is the probability of having been admitted to the hospital for diabetes related acute conditions (diabetic ketoacidosis or hypoglycemic coma). Mean differences relative to the comparison group are reported with 95% CI. The outcome mean is 0.37.

Figure 9: Difference in probability of diabetes related complications by parental characteristics

Panel A: Difference by paternal education

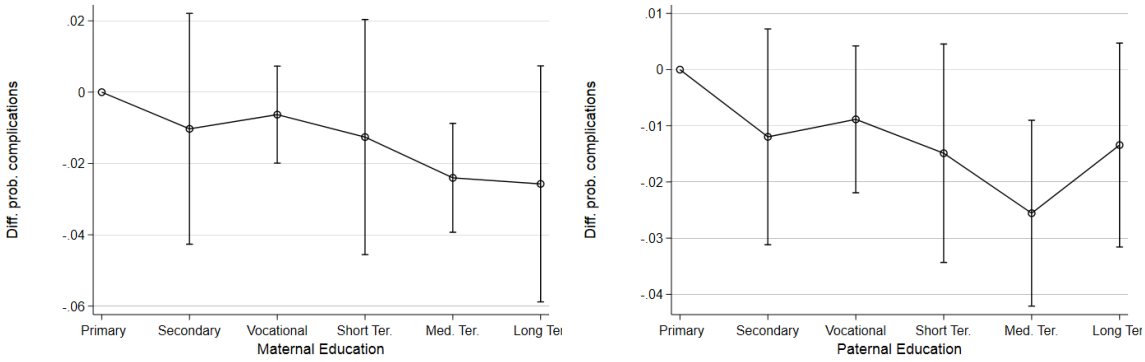

Panel B: Difference by paternal income quartile

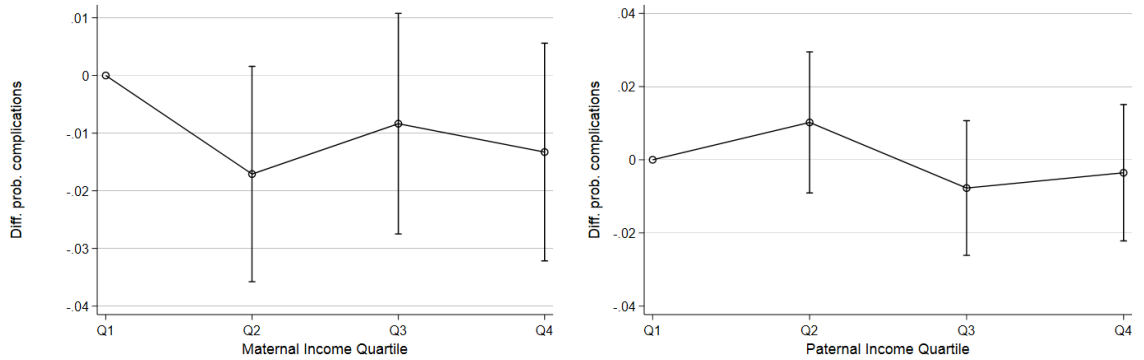

Panel C: Difference by parental age

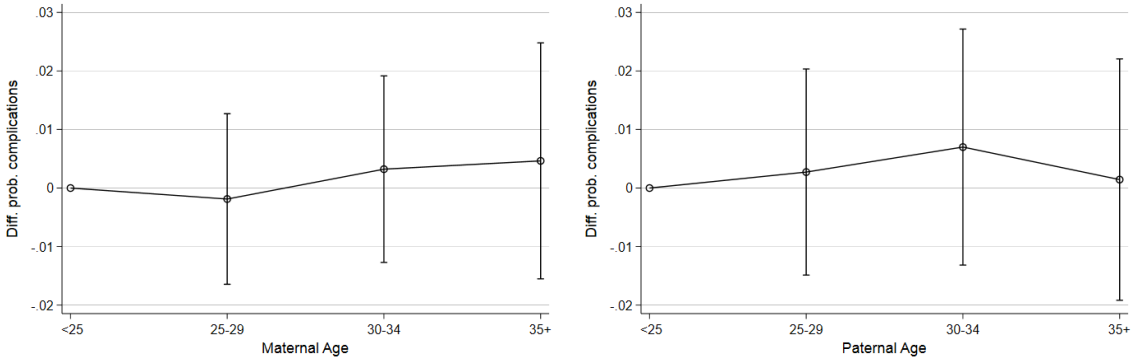

Notes: This figure presents subgroup differences in diabetes related outcomes among the individuals with diabetes. The outcome is the probability of having been diagnosed with any late complication by age 30. Mean differences relative to the comparison group are reported with 95% CI. The outcome mean is 0.54.