Intergenerational Spillovers in Disability Insurance*

Gordon B. Dahl† Anne C. Gielen‡

July 25, 2017

NOTE: Preliminary and incomplete draft

Abstract: Does participation in a social program by a parent have spillovers on their child's use of public assistance, future labor market outcomes, and human capital investments? From a policy perspective, what a child learns from his or her parents about employment relative to government support could matter for the financial stability of a variety of social insurance and safety net programs. While intergenerational spillover concerns have figured prominently in policy debates for decades, the evidence base is scarce due to nonrandom participation and data limitations. In this paper we exploit a policy reform in the Netherlands which tightened disability insurance (DI) criteria for new applicants and existing claimants, and use rich panel data to link parents to children's long-run outcomes. The key to our regression discontinuity design is that the reform applied to younger cohorts, while older cohorts were exempted from the new rules. We find that children of parents who were pushed out of DI or had their benefits reduced are 11 percent less likely to participate in DI themselves and earn two percent more in the labor market as adults. There is no effect on children's participation in other government safety net programs. As a result, intergenerational spillovers caused both reduced government transfers and increased tax revenue, resulting in 5,600 euros in positive child spillovers per treated parent. Children anticipate a future with less reliance on DI by investing in extra education, with treated children completing an extra .12 years of education on average. Our findings have important implications for the evaluation of the costs and benefits of this and other policy reforms; indeed, ignoring parent-to-child spillovers underestimates the cost savings of the Dutch DI reform in the long run by roughly 20 percent.

Keywords: Peer effects, disability insurance, intergenerational links
JEL codes: I38, H53, J62

---

*We thank Kate Antonovics, Prashant Bharadwaj, Julie Cullen, Roger Gordon, and Erik Plug for helpful advice, and seminar participants at several universities and conferences for helpful comments and suggestions. Financial support from Erasmus University Rotterdam is gratefully acknowledged.

†Department of Economics, University of California San Diego; email: gdahl@ucsd.edu
‡Erasmus School of Economics, Erasmus University Rotterdam; email: Gielen@ese.eur.nl

1 Introduction

Does participation in a social program by a parent influence their child's own social program participation, future employment and earnings, and human capital investments? From a policy perspective, what a child learns from his or her parents about employment relative to government support could matter for the financial stability of a variety of social insurance and safety net programs. Indeed, arguments about the presence, type and size of intergenerational spillovers have figured prominently in policy debates for decades. Some policymakers claim that parental participation creates a cycle of government dependence and reduced employment, while others believe intergenerational patterns simply reflect shared, negative environmental factors.

Evaluating the intergenerational effects of income support programs is difficult because a parent's participation is not randomly assigned. On the one hand, observing a parent out of the labor force and on public assistance could alter a child's perceptions about the relative costs, benefits, and stigma associated with the two alternatives. There could also be information transmission or differential child investments as a result of having a parent receive government transfers. On the other hand, characteristics like poor health, bad neighborhoods, or reduced employment opportunities could be correlated across generations, creating mechanical intergenerational links which do not reflect a behavioral response of the child.

Identifying intergenerational spillovers in program participation is a difficult empirical problem. Simple correlations are unlikely to capture a causal effect, as unobservable characteristics which are correlated across generations or influence a family's environment are likely to bias the estimates. Credible identification requires an exogenous shock which affects a parent's participation, but does not directly affect their children. On top of this, the researcher needs a panel dataset which links parents to children, contains a detailed set of outcome variables, and follows families over a long period of time. This paper overcomes these challenges using a policy reform which generated quasi-experimental variation and a rich administrative dataset.

Our setting is disability insurance (DI) in the Netherlands and a 1993 reform which was prompted by the rising costs of the Dutch system. In 1969, two years after its introduction, 4% of the Dutch working age population participated in disability insurance, but by the late 1980s, participation had risen to 12%. At its peak, the program cost 4.2 percent of GDP, and was not fiscally sustainable. Similar trends, while not as dramatic, have occurred in most industrialized nations, including the U.S., the U.K., and other countries in Europe (see Burkhauser, Daly, McVicar, and Wilkins, 2013). Due to a series of reforms, including the one we study, Dutch DI payments now constitute around

2.1% of GDP, which compares to 2.3% on average in Europe and 1.7% in the U.S.

The 1993 reform we focus on simultaneously tightened eligibility criteria and lowered payment generosity. It forced current DI recipients to be re-examined by a medical doctor and subjected to a new set of rules which unambiguously made them (weakly) worse off. Some individuals received lower payments and others were disqualified from the program entirely. Importantly, the more stringent re-examination rules only applied to individuals less than age 45 as of August 1, 1993, since at the last minute it was decided that individuals age 45 or older would be grandfathered in under the old program rules. This differential application of the new rules creates an age discontinuity in program eligibility; individuals around the cutoff should be similar in all dimensions except for the increased stringency of DI re-examinations.

Using a regression discontinuity (RD) design, Borghans, Gielen, and Luttmer (2014) find that approximately 4 percent of DI participants exited DI due to the more stringent rules and that annual benefits fell by around 1,000 euro, or 10 percent. We find slightly larger effects for our sample of parents with children still living at home at the time of the reform, with 5 percent of parents exiting DI due to the reform and annual benefits dropping by 1,300 euro on average. Borghans et al.'s analysis also reveals a strong rebound in labor earnings of 0.62 euros on average per euro of lost DI benefits and a .30 euro increase from other social assistance programs in the short run.

The goal of our paper is to explore how a parent's reduction in DI benefits, with its resulting changes in employment and social support substitution, affects their children's choices. We focus on children's future participation in DI and other social programs, labor market outcomes as adults, and human capital investments when young. It is important to note the DI rule changes affected parents on both the intensive and extensive margins: some parents had their DI payments reduced while others were kicked off the program entirely. We therefore focus on the reduced form effects of the DI reform on child outcomes, but also present IV estimates which scale these effects by the parental drop in DI payments (treating exit from the program as a reduction in payments to 0). We use an RD design, where the running variable is the age of the parent and the dependent variables are various outcomes for their children.

Our first result is that there exists a significant causal link in DI usage between parents and children when they become adults. Children whose parents exit DI or have their benefits reduced are 1.1 percentage points less likely to have ever participated in DI 21 years after the reform date (in 2014, when they are 36 years old on average). This is a sizable 11 percent drop relative to the mean. Zooming in on cumulative usage of the program, children are estimated to be on DI 47 fewer days by 2014 (including zeros) if they had a parent subject to the new rules, which represents a 16 percent drop. Using cumulative income received from DI as the dependent variable instead, children received roughly 1,600 euros less in DI payments (including zeros), which is sizable relative to the overall mean of 10,100 euros in cumulative DI receipt for these children.

To get a fuller picture of intergenerational spillovers and fiscal impacts, we next assess whether a child's taxable earnings and participation in other social support programs change. Only with this information can one calculate the total spillover effect on the government's budget, which includes all transfers and taxes. We find that cumulative earnings rise by approximately 7,200 euros (a little less than 2%) for children of parents subject to the less generous DI rules. In contrast, we find no detectable change in cumulative unemployment insurance receipt, general assistance (i.e., traditional cash welfare), or other miscellaneous transfer programs. The estimated cumulative increase in taxes minus government transfers is approximately 3,300 euros per child. While 1,400 euros of this amount is due to cost savings from lower DI payments, the remaining 1,900 is due to increased tax revenue resulting from higher earnings. Since parents in our sample have an average of 1.7 children living with them at the time of the reform, this implies 5,600 euros in positive intergenerational spillovers per treated parent.

We then turn to children's educational investments. When a parent is forced off of DI or has their benefits reduced, their child invests in a statistically and economically significant 0.12 extra years of education. The largest increase occurs for the margin of graduation from upper secondary school (the equivalent of high school in the U.S.). Since most schooling takes place before children have entered the labor market, these findings provide intriguing evidence of an anticipatory effect: children whose parents experience a reduction in DI benefits plan for a future with less reliance on DI in part by investing in their labor market skills. Indeed, the schooling effect is concentrated on children who are less than age 18 at the time the reform was implemented.

Finally, we explore two other outcomes which could be affected: crime and marriage. We find a modest reduction in future incarceration, but no significant effect on arrests. There is also some evidence that children are more likely to get married if their parent was exposed to the new, less generous DI rules.

Taken together, these results indicate that children respond strongly when a parent exits DI or has their benefits reduced. Understanding the Dutch context is key for interpreting these intergenerational effects. As a result of the reform, parental leisure decreased and work hours increased substantially, with total parental income changing little in the short run but declining in the long run. In theory, less parental supervision due to increased work hours or lower income in the long run could result in reduced investments in children, harming their attachment to the labor force as adults. Based on our estimates, however, this is not the case. Instead, the findings all point to a greater focus by children on formal employment. Children increase their educational investments, earn more in the labor market, and decrease their DI use. This is consistent with forward-looking children anticipating they will rely relatively less on government assistance in the future.

The reform we study affects parents who are all on DI to begin with. Therefore, the spillover effects we observe cannot be due to information about how to apply for DI in general, or what it is like to be on DI. Instead, children of affected parents are learning about the transition to work or other social programs after a reduction in DI support. Interestingly, we find no evidence that children participate more in other social support programs, like unemployment insurance, even though their parent's use of these programs went up in the short run.

While intergenerational spillover effects of program participation are central for policy discussions, there is remarkably little existing causal evidence. There are several observational studies which document intergenerational links, but few with credible research designs (see Black and Devereux 2011 for a survey). Some studies attempt to solve the endogeneity problem using either a rich set of covariates or family fixed effects (e.g., Solon et al. 1988, Gottschalk 1996). Four papers that we know of have tried to come up with exogeneous sources of variation in program participation for identification.[1] Antel (1992) uses state-level welfare benefits and net migration flows as exclusion restrictions in a Heckman selection model and finds evidence for intergenerational links. Levine and Zimmerman (1996) uses variation in state benefit levels and local labor market conditions as instrumental variables and concludes that most of the intergenerational correlation in welfare use is not causal. Hartley, Lamarche and Ziliak (2017) uses variation across U.S. states in the timing of welfare reform implementation and finds a mother's use of welfare significantly increases the chances her daughter will participate as well. Finally, Dahl, Kostol, and Mogstad (2014) uses a random judge assignment design and finds that DI participation by parents in Norway increases the chances their child will participate as well.

Since the paper by Dahl et al. is closest to ours in terms of also studying intergenerational DI participation, we briefly contrast it with the current paper. As background, their sample is comprised of the 6.25% of all DI applicants who are initially denied but then appeal, since this is the group which had random assignment of judges who differ in their stringency.[2] Relative to their paper, we make several contributions. First, we examine the effects of an actual policy reform which

[1] In a related paper, Pepper (2000) illustrates the difficulty in drawing credible inference using observational data or standard instruments, such as local unemployment rates, for studying intergenerational welfare receipt.

[2] The treatment in Dahl et al. is a combination of some parents never being allowed on DI as well some parents having delayed entry into DI; 75% of parents who are denied at the appeals stage eventually reapply, with 65% of these being ultimately allowed DI.

was implemented nationwide and had large fiscal consequences. Second, we can precisely estimate a broader set of labor market, public assistance, and human capital investment outcomes. This is mostly because Dahl et al. cannot estimate spillover effects on education or employment with enough precision to draw firm conclusions; our estimation sample is 8 times larger and follows children for 21 years after the reform date. Third, we are able to calculate the total fiscal costs to the government budget due to intergenerational spillovers, including changes in all taxes and transfers, and not just DI payments.[3] Fourth, we are able to follow children for a longer time period, when DI use is more common among the children.[4] Fifth, we examine the consequences of pushing individuals off of DI or reducing their benefits, whereas Dahl et al. study the effects of making it harder to get onto DI in the first place. There is no a priori reason to expect intergenerational effects will be symmetric for the entry and exit margins, and any asymmetry matters when considering specific public policy reforms. And finally, we use a completely different, but equally compelling, quasi-experimental research design to better understand this important, but largely unanswered, question.

More broadly, our study complements a related literature which looks at other shocks to parents which have the potential to change children's long-run outcomes. Oreopoulos, Page and Stevens (2008) find that after a parental job loss due to a firm closure, children have lower earnings and higher UI and social assistance participation. Likewise, Rege, Telle, and Votruba (2011) find that children perform worse in school after a job loss due to a plant closure, particularly if the father was affected. Stevens and Schaller (2011) find similar effects using observational data on job loss. A series of papers look at the effects of income shocks, such as those due to child tax benefits (Dahl and Lochner, 2012; Milligan and Stabile, 2011) or DI changes (Chen, Osberg, and Phipps, 2015) and generally find increases in academic achievement and child well-being. A different type of shock is having a parent move to a better neighborhood. Initial analyses of adults and older youth from the Moving to Opportunity (MTO) experiment found no effect on earnings or employment (e.g., Katz, Kling, and Liebman 2001; Oreopoulos 2003), while subsequent work by Chetty, Hendren, and Katz (2016) find large effects from this experiment for young children at the time of the move, and conclude that better neighborhoods have the potential to reduce the intergenerational persistence of poverty.[5]

[3]This is empirically relevant, as the net fiscal impact is more than twice as large as the intergenerational savings based on reduced DI payments alone. A similar point, although not in the intergenerational context, has been made in a recent working paper by Autor, Kostol, Mogstad, and Seltzer (2017).

[4]In Dahl et al.'s baseline sample, only 3 percent of children have participated in DI. In contrast, 10 percent of the children in our baseline sample have participated in DI.

[5]There is also a related literature on disability insurance programs and their labor supply effects. See Autor and Duggan (2006), Autor (2011), Bound (1989), Bound and Burkhauser (1999), Chen and van der Klaauw (2008), Campolieti and Riddell (2012), De Jong, Lindeboom, and van der Klaauw (2011), French and Song (2014), Gruber (2000), Gruber and Kubik (1997), Kostol and Mogstad (2014), Maestas, Mullen, Singleton (2012), Strand (2013), and

The remainder of the paper proceeds as follows. The next section provides background on disability insurance in the Netherlands, the 1993 reform, and the data. Section 3 lays out our RD design and discusses threats to identification. In Section 4, we present the first stage estimates for parents. Sections 5 and 6 present our main results on child spillovers in program participation, work, and education. Section 7 conducts some heterogeneity and robustness analysis. Section 8 discusses and interprets our results, followed by a short conclusion in Section 9.

2 Background and Data

2.1 Disability Insurance in the Netherlands

The modern Dutch DI program was created in 1967 by merging two existing programs covering workplace-induced injuries and disabilities unrelated to employment. The program was generous compared to other countries, as it covered all workers with no waiting period, replaced up to 80 percent of wages, and included a variety of subjective illnesses. Moreover, sickness benefits replaced a worker's wages between 80 and 100% during the transition to disability insurance, and workers on sickness benefits for a full year were routinely transferred to the DI program without a serious reappraisal of their disability (Kalwij, de Vos, and Kapteyn, 2014). These factors fueled a rapid rise in DI recipients, from 4 percent participation of the eligible population in 1967 to over 8 percent by 1980. Modest reforms in the early 1980s were enacted in an attempt to stem the rise, but were largely ineffective. Participation reached a peak of 12% of eligibles in the late 1980s, with payments ballooning to 4.2% of gross domestic product.

Starting in the 1990s, a series of reforms were implemented to control the spiraling costs of the DI system, including reductions in benefit levels, tightened eligibility criteria, changes to the sickness benefit program, and increased financing and responsibility transferred to individual employers. The cumulative effect of these reforms was that by 2012 the participation rate had fallen to just over 7% of the eligible population. Going forward, the participation rate is predicted to fall even further as the stock of older recipients transitions out of the DI program and on to the retirement pension program. The trends over time are documented in Figure 1 and discussed in more detail by Koning and Lindeboom (2015).

The current state of DI in the Netherlands is that payments now total around 2.1% of GDP. This compares to 2.3% in other European countries, and 1.7% in the U.S. In terms of participation, the

von Wachter, Song, and Manchester (2011). Most of these studies find sizable labor supply responses to DI benefit generosity.

2012 Dutch rate of 7% is higher compared the U.S. rate of 5%, but lower than Norway's 10%, for example. One interesting contrast is that the U.S. rate continues to rise and is projected to reach 7% by 2018 (Burkhauser and Daly, 2012), while the Dutch rate is continuing to fall. Because of this, several policy analysts have proposed adopting several aspects of the Dutch system (as well as the German and U.K. systems) to reverse the steeply increasing DI trends in the U.S.

Before continuing, we note several differences between the Dutch and U.S. programs. First, in the Netherlands, individuals can receive payments for a partial disability and therefore continue to work and earn benefits simultaneously, while in the U.S. disability determination is binary. Second, health insurance and other benefits are unrelated to DI receipt in the Netherlands, but directly linked in the U.S. Third, benefits do not depend on family size in the Netherlands, while they do in the U.S. Fourth, the replacement rate in the Netherlands is not a function of tenure, with all workers being covered 100% the first day on the job. Finally, the replacement rate of 70% for complete disability in the Netherlands is higher than the average U.S rate of 40 to 50% (see Borghans et al., 2014; Duggan, 2003).

2.2 1993 Reform

Many changes are responsible for the reduction in DI expenditures in the Netherlands; in this paper we take advantage of a 1993 reform which generates a discontinuity in program generosity based on age. As this is the same cohort discontinuity used by Borghans, Gielen, and Luttmer (2014) to study benefit substitution, we only briefly explain the most salient features of DI in the Netherlands and the 1993 reform, and refer readers to their paper for further details.

In the Netherlands, individuals receive DI payments based on the degree of their disability, which is based on the calculated income loss due to a disability. Calculated income loss is determined by comparing pre-disability earnings to a constructed measure of "earnings capacity." The reform that we exploit affected the calculation of this "earnings capacity," making it less generous to DI claimants.

The degree of disability is denoted in 8 categories; which category an individual belongs to is determined by the ratio of pre-disability earnings minus earnings capacity to pre-disability earnings. Individuals can continue to work and earn up to their remaining earnings capacity (pre-disability earnings minus earnings capacity) after a disability, and at the same time receive DI payments for the fraction of lost earnings.[6]

To explain the cohort discontinuity, we first need to describe how earnings capacity and benefits

[6] Pre-disability annual earnings are indexed and subject to a cap (roughly 36,000 euro in 1999). If individuals earn more than their capped earnings exemption, their DI benefits are reduced temporarily, with a reclassification of the degree of disability only happening if an individual exceeds the cap for three years.

were determined before and after the 1993 reform. Prior to 1993, a medical doctor examined applicants and created a subjective list of work activities the applicant could still perform, based on a set of 27 physical activities (e.g., lifting, kneeling) and 10 psychological abilities (e.g., the ability to work under time pressure). This work activity list, in conjunction with the applicant's education level, was used to create a list of suitable occupations from a dictionary of occupational requirements. The applicant's earnings capacity was then defined as the average wage in the 5 highest-paying suitable occupations which had at least 10 active workers in the applicant's geographic region.[7] If 5 suitable occupations could not be found, earnings capacity was set to 0. The calculated degree of disability was then binned into categories which determined the replacement rate. Replacement rates varied from 0 to 70 percent of prior earnings.[8]

The 1993 reform altered this process in two ways. First, it mandated the doctor create a list of work activities based on a more objective medical diagnosis which could be directly linked to functional work limitations. Second, (i) the list of suitable occupations was expanded by no longer taking education level into account, (ii) only 3 suitable occupations were used to calculate earnings capacity, and (iii) the geographic region of 10 active workers was expanded to be three times larger. Each of these changes weakly reduce the degree of disability for an applicant compared to the old criteria, as remaining earnings capacity can only rise. Moreover, the new rules make it more likely that enough suitable occupations can be found, reducing the chances of total disability. The end result is that fewer individuals qualify for DI and benefit levels are weakly reduced for those who continue to qualify, as long their disability has not gotten worse since their last re-examination.

A feature of the 1993 reform is that it specified all individuals age 50 or older at the time of the reform would be subject to the old rules and not re-examined at all. For individuals below age 50 as of January 8, 1993, the new rules affected both new applicants and existing DI participants. Since it was not logistically feasible to re-examine all DI participants immediately, they were scheduled to be re-examined over the ensuing years based on their age cohort, starting with the youngest cohorts under the age of 35 on August 1, 1993. The 35 to 40 year old cohort was scheduled to be re-examined in 1995, 41-45 year old cohort between 1996-1997, and the 45-50 year old cohort between 1997-2001. However, on November 12, 1996 the Dutch Parliament passed a motion grandfathering the 45-50 year age group into the old, more generous rules. This grandfathering creates a sharp cutoff in the generosity of DI based on an individual's age, a feature we exploit for identification.

[7] If there are not initially 5 suitable occupations within an applicants narrow region, the region is expanded.

[8] For a degree of disability between 80-100% the replacement rate is 70%, for 65-80% it is 50.75%, for 55-65% it is 42%, for 45-55% it is 35%, for 35-45% it is 28%, for 25-35% it is 21%, for 15-25% it is 14%, and for less than 15% it is 0%.

2.3 Data

Our analysis uses several data sources that we can link through a unique identifier assigned to all individuals in the Netherlands. We combine administrative data from several sources on the universe of children of DI recipients for the time period we study. The disability administrative records begin in 1996 and are observed as late as 2014. The records include information on the start and end dates of a spell, the binned disability rating, DI payments received, pre-disability earnings, and the reason a spell ends. The records do not contain the medical doctor's diagnosis, the list of work activities the individual could still perform, or the set of suitable occupations.

We merge in data from a variety of administrative records for the period 1999 to 2014. We use data from Statistics Netherlands for earnings, self-employment, and unemployment insurance which is compiled using information from three different tax and social insurance record sources. This data starts in 1999. Unemployment insurance in the Netherlands can last up to 5 years depending on prior work history.

Data on general assistance (traditional cash welfare) and miscellaneous benefit programs come from the various organizations that administer the programs. As opposed to the U.S., general assistance has no time limit in the Netherlands and does not require dependents, although it is means tested. There are about 30 miscellaneous benefit programs, most of which are small in terms of benefit amounts and the size of the eligible population.

We further merge in educational attainment as of 2014, as well as family structure in 2014. The education data is complete for younger cohorts, but only a sample for the older cohorts. Crime data on arrests and incarcerations come from two different data sources, and both span 2005-2014. Finally, we use municipal registry files for basic demographics. One advantage of this rich dataset merged from several sources is that we can study a variety of spillover effects across generations. Further details on most of these variables, and how they are measured, can be found in Appendix B of Borghans, Gielen, and Luttmer (2014).

We limit our sample to children of parents who were between the ages of 40 and 50 as of the reform date of August 1, 1993. We also require the child to still be living at home on August 1, 1993; as an extra specification, we estimate effects for children not living at home at the time of the reform. Parents with multiple children appear more than once in the sample, while children with two parents on DI are dropped from the sample. Due to data availability, our sample is limited to parents who were receiving DI benefits on August 1, 1993 and who were were still on DI as of January 1, 1996 (the earliest date for which we have DI records). It is important to realize that even though we do not have DI records before 1996, this should not create any biases. The reason is that 1996 is still before the DI re-examinations took place for the age 40-45 cohort and before the passage of the DI rule change exempting the age 45-50 cohort. Starting with 1996 affects the interpretation of our estimates, but not their validity. After imposing all of these restrictions, we have a sample of 101,125 child observations. For the education analyses, our sample is smaller (N=65,208) since education was collected for all individuals in later cohorts, but only a subsample of earlier cohorts.

Summary statistics for both parents and children can be found in Appendix Table A1. The first column displays sample means for parents who were between the ages of 40 to 50 as of the cutoff date, and still on DI as of 1996. The other two columns show means for subsamples on each side of the 45 year-old age cutoff. On average, a parent has been on DI for a little over 6 years as of the reform cutoff date, with the older sample having an 6 more months of participation. Fifty-eight percent of parents are classified as fully disabled. Older parents are 10 percentage points more likely to be fully disabled, while younger parents have higher rates of low-level disability. Parents in our sample are predominantly male, married, and native Dutch.

We have 70,319 parents matched to 116,356 children who still live at home as of the reform date of August 1, 1993. The average age of children as of the reform date is 15 years old, although it should be remembered that re-examinations of the 40 to 45 year olds did not start until 1996 and that the 45 to 50 age group was not grandfathered in under the old rules until the end of 1996. By this time, children are more than 3 years older (roughly 18 years old on average). Children in our sample are somewhat more likely to be male, reflecting the fact that sons tend to move out of their parent's house at an earlier age.

3 Model

3.1 Regression Discontinuity Design

The discontinuity we exploit arises from the fact that the 1993 reform affected some DI participants, but not others, based on their age. Parents who were age 45 to 50 as of August 1, 1993 were subject to the old DI rules, while parents between the the ages of 40 to 45 were re-examined according to the new, more stringent rules. The direct effect of the reform on parental outcome y^P can be modeled in a regression discontinuity (RD) framework as:[9]

$$y_i^P = \alpha^P + 1[t_i \geq c](g_l(t_i - c) + \theta) + 1[t_i < c]g_r(c - t_i) + \delta^P x_i + e_i^P \tag{1}$$

[9]See Imbens and Lemieux (2008) and Lee and Lemieux (2010) for details on the implementation and assessment of RD designs.

where t is the age of the parent on August 1, 1993, c is the cut-off age of 45, x is a vector of pre-determined parental and child characteristics, e^P is an error term, and g_l, and g_r are unknown functions. The coefficient θ is the first stage coefficient for the associated parental outcome (DI payment amount, or alternatively, DI participation).

The reduced form model for our RD design can be implemented as:

$$y_i^C = \alpha^C + 1[t_i \geq c](h_l(t_i - c) + \lambda) + 1[t_i < c]h_r(c - t_i) + \delta^C x_i + e_i^C \tag{2}$$

where y^C is the relevant child outcome variable, t is the age of the parent on August 1, 1993, c is the parental cut-off age of 45, x is a vector of pre-determined parental and child characteristics, e^C is an error term, and h_l, and hg are unknown functions. The coefficient λ is the reduced form (RF) or intention to treat (ITT) effect of the reform on outcomes. In the absence of covariates, the IV estimate is simply the ratio of the RF estimate of λ to the relevant first stage estimate of θ.

3.2 Threats to Identification

3.2.1 Manipulation

The validity of an RD design requires that individuals cannot manipulate the assignment variable, which in our setting is the parent's age at the time of the reform. Since parents cannot change their actual or officially recorded age easily in the Netherlands, there is little chance for this type of direct manipulation.

Since the DI data is not available until 1996, another threat to validity is that the reform caused differential attrition around the age 45 cutoff. As a reminder, our sample includes parents who were receiving DI benefits on August 1, 1993 (the date the reform went into effect for new applicants and the youngest cohort of existing claimants) and who were were still on DI as of January 1, 1996. In other words, we can only observe whether an individual was receiving DI at the time of the initial implementation of the reform if they remained on DI until at least January 1, 1996. While the reform likely caused some claimants to exit DI in anticipation that they would be re-examined, it is unlikely to have caused a jump in exits around the age 45 cutoff. The reason is the re-examinations for individuals age 40-45 did not start until after January 1, 1996 and it was not until November 1996 that Parliament decided the 45-50year old cohort would be grandfathered in to the old, more generous rules.[10]

[10] While 40 years old were initially scheduled to be re-examined at the end of 1995, the re-examinations took longer than initially expected. In conversations with the disability insurance office, we learned that few of the 40 year old cohorts were re-examined until after 1996.

Borghans et al. (2014) perform two empirical tests for manipulation for their sample, which includes all individuals on DI, and not just parents. They first graph the histogram of age at the time of the reform, and find no noticeable jumps around the age 45 cutoff. We find a similar result for our sample of parents: using a McCrary (2008) test, we do not reject the null hypothesis of a smooth density around the 45 year old cutoff (p-value=.33). Second, they find no systematic evidence of changes in the distribution of pre-determined characteristics around the reform date. Using our sample of parents, we similarly find that almost all of the pre-determined characteristics do not jump significantly at the 45 year old cutoff. Moreover, the point estimates are small in magnitude and our RD estimates barely move when we include these characteristics in the regressions.

3.2.2 Exclusion Restriction

As long as parents cannot manipulate their age and there is no differential attrition around the age cutoff, the RD design will identify the ITT effects for children. That is, we can estimate the causal impacts on children of the 1993 DI reform which tightened DI generosity for some parents but not others. To scale these reduced form effects, we will be using parental DI payments as the first stage outcome. Interpreting the resulting IV estimates as the causal effect of a drop in parental DI payments requires an exclusion restriction: whether a child's parent was exposed to the 1993 reform should affect the child's outcomes only through the drop in parental DI payments, and not directly in any other way.

It is unlikely parental exposure to the 1993 reform affected children directly except through the reduced generosity of the DI program. However, the drop in DI payments may not be a sufficient statistic for how generosity changed at the cutoff. For parents remaining on the program, the reform (weakly) decreased DI payments, whereas for parents kicked off DI or choosing to leave voluntarily, the reform reduced their payment to zero. Parental DI payments will capture both the intensive and extensive margins of the reform under the assumption that total DI payments are what matters. For the exclusion restriction to hold, therefore, parental participation versus non-participation cannot directly affect children except through the reduction in payments to zero. This implies, for example, that a reduction in benefits from 10,000 to 7,000 euros has the same effect as a parent who previously received 3,000 euros exiting the program and receiving 0 euros. Since this may not be the correct functional form for how the new stricter rules affected children, we focus more on the reduced form estimates throughout.

The 1993 reform may also have triggered a variety of changes for exposed parents, such as changes in parental labor supply, available family income, or even family structure. It is important to note these changes do not violate the exclusion restriction. Instead, they are potential mechanisms through which a shock to parental DI generosity affects children.

3.2.3 Monotonicity

If the effect of the drop in parental DI payments is constant for each child outcome, then the absence of manipulation combined with the exclusion restriction are sufficient for consistent IV estimation. With heterogeneous effects, however, monotonicity is also needed. In our setting, monotonicity requires that if a parent was exposed to the new, more stringent DI rules, they must receive DI payments which are lower or the same compared to what they would have received under the old rules. Monotonicity ensures that IV identifies the local average treatment effect (LATE) of a drop in parental DI payments, that is, the average effect among the subgroup of children whose parent's DI payments would have been lowered if they were exposed to the new versus old rules.

Since the new rules weakly reduced payments for any individual whose situation has not changed, monotonicity holds by construction for most of the sample. The one exception is that if a parent's illness has worsened, the required re-examination under the new, stricter rules could result in a higher degree of disability classification (and hence a higher DI payment). The potential worry is that some parents whose condition has worsened would not have otherwise requested a re-examination and been reclassified. Comparing the 40-45 age cohort, which had mandatory re-examinations, with the 45-50age cohort, which were not required to be re-examined, reveals this is unlikely to be an important issue. For the 40-45 age cohort, 5.8 percent of the sample had their degree of disability rating increase between 1996 and 1999 (the period of the mandatory re-examinations), whereas for the 45-50age cohort, 6.6 percent had their rating increase. This comparison reveals that any margin for non-monotonicity to matter is small, even taking into account that rating increases are expected to occur somewhat more often for older individuals.

4 First Stage Parental Estimates

This section documents the effect of the reform on parents using an RD design. An advantage of RD is that results can be presented graphically, which provides a transparent way of showing how the intergenerational spillovers are identified. Throughout the paper, we will begin with a graphical depiction of key outcomes before turning to a more detailed regression-based analysis. The figures will include outcomes aggregated into parental age bins, as well as separate linear trends on each side of the cutoff estimated using the underlying data and baseline regression specification. The regression lines best illustrate the trends in the data and the size of the jump at the age cutoff, whereas the binned means provide a sense of the underlying variability in the data.

The top panel of Figure 2 graphs the relationship between parental DI payments and the reform. The sample includes parents who were receiving DI benefits before the reform, as detailed in Section 2.3. The running variable is the parent's age as of the reform date of August 1, 1993 and the cutoff date of 45 years old determines whether the parent is subject to the new versus old DI program rules. On the y-axis is parental DI benefits in 1999; we use 1999 since this is after all the re-examinations have taken place. Our age variable is recorded at the monthly level; each observation in the graph is the average DI payment for parents in six-month age bins. Three-month age bins for this graph, and all other RD graphs, can be found in the Appendix.

The figure reveals that DI benefit payments rise with age, largely reflecting the fact that older individuals have higher degree of disability ratings on average and therefore higher DI payments. More importantly, there is a sharp drop in payments for individuals just to the left of the cutoff. This is as expected, since parents less than age 45 were subject to re-examination under the stricter DI program rules. DI payments drop by around 1,300 euros, which is a reduction of 10 percent compared to the average benefit level in the sample.

To document the extensive margin of the DI reform by itself, in the bottom panel we graph the fraction of parents who exit DI completely. The running variable and cutoff are the same as in the top panel. Each observation in the graph is the fraction of parents in a six-month bin who have exited DI by 1999. The first pattern to notice is that exits are higher for younger individuals, with over 20% of 40 year olds exiting compared to around 5% of 50 year olds. At the cutoff, there is a sizable 5 percentage point drop in exits. This drop reflects individuals being kicked off of DI under the stricter re-examination rules, but could also include any voluntary exits.

In Table 1 we present regression results corresponding to these figures. Our baseline specification, here and in what follows, regresses the relevant outcome on a dummy for the reform cutoff and separate linear trends in parental age to the left and the right of the cutoff. We use triangular weights so that observations nearer the cutoff will have more influence. Although the coefficients are not reported, we also include a variety of covariates for both the parent and the child which are measured as of 1996.[11] For the parent, we include age, birth month dummies, a gender dummy, a cubic in pre-disability earnings, six dummies for degree of disability, a cubic in DI duration, and

[11] The year 1996 is before the passage of the law exempting the 45-50 age cohort from the new, less generous DI rules and before the re-examinations have occurred for the 40-45 age cohort, so these controls should be exogenous to the cutoff. Note also that 1999 is after all re-examinations have taken place and that all new applicants were subject to the stricter guidelines, regardless of age.

national origin dummies; for the child we include age, birth month dummies, and a gender dummy.

The first specification in Table 1 uses DI benefits received by parents in the year 1999 as the outcome. Mirroring what was drawn in the top panel of Figure 2, the first stage RD estimate is sizable a 1,300 euro drop in benefits for parents exposed to the reform. This first stage point estimate is more than 13 times its standard error. Both the size and the precision of this estimate are important for identifying spillover effects on children, which by their nature are second order effects. The second specification uses exit from DI by 1999 as the outcome, and finds a large and precisely estimated 5.4 percentage point drop in exits at the age cutoff.

As a reminder, the reform mandated the 40-45 age group be re-examined and subjected to rules which lowered the calculated degree of disability and therefore the replacement rate. This meant that while some individuals exposed to the reform were kicked off the program, others remained on DI but with lower benefits. Given the reform had both an extensive and intensive margin, we focus primarily on the reduced form, or intent to treat, estimates when analyzing children's outcomes. But to provide a sense of scale, we also use the total drop in DI payments (including drops to zero) for parents as a first stage to construct an IV estimate. Since exit is likely to have a direct affect beyond the reduction in DI benefits to zero, the exclusion restriction necessary for IV is unlikely to hold. For this reason, we present both reduced form and IV estimates: the RF estimates will be valid regardless, while the IV estimates should merely be viewed as providing a possible scaling of the effects.

As a results of the reform, other parental outcomes causally changed as well, including increased employment and substitution to other government benefit programs. We discuss this when we interpret our findings in Section 8.

5 Spillovers in Program Participation and Work

5.1 Child DI Participation

We begin our investigation of intergenerational spillovers by exploring the linkage in DI participation between parents and their children. Figure 3 presents RD graphs for the extensive and intensive margins of DI use. The x-axes in both graphs are the same as in Figure 2, with the running variable being the age of the parent as of the reform date and the cutoff age of 45 being marked with a vertical line. But now the y-axis plots the child's participation in DI, rather than the parent's. The graph also includes separate linear trend lines on each side of the cutoff date along with pointwise 95%

confidence intervals. These trend lines are based on the underlying, unaggregated data for a parent's age. As before, each observation in the graph is an average for six-month age bins; three-month age bins can be found in the Appendix.

A key advantage of our long panel is that we can measure outcomes when the children are much older, after they have had a chance to live on their own, enter the labor market, and participate in the DI program. For our main child outcomes, we measure cumulative effects as of 2014, which is 21 years after the reform which affected their parents was announced. By this time, children are 36 years old on average, with the range of child ages spanning from 28 years old at the 10th percentile to 40 years old at the 90th percentile. Between 1999 and 2014, over 10 percent of children in our sample have participated in DI at some point, with an average number of 298 days spent on the program (including zeros).

The top graph in Figure 3 looks at whether a child has ever participated in DI between 1999 and 2014 as the outcome. There is a noticeable jump in child DI participation at the parental age cutoff of 45. Likewise, there is a noticeable jump in the cumulative number of days a child has been on DI. Table 2 presents the reduced form estimates corresponding to these graphs. For the extensive margin of participation, there is a statistically and economically significant 1.1 percentage point drop for children if their parent was exposed to the reform. This is an 11% effect relative to the mean. Likewise, children participate in DI for 47 fewer days if their parent was subject to the stricter DI rules, which represents a 16% drop relative to the mean.

While the reduced form estimates are preferable for reasons discussed earlier, the table also presents IV estimates to provide a sense of scaling. For these IV estimates, we use the total drop in parental DI payments, including drops to zero, as the first stage outcome variable (i.e., panel A in Table 1). These IV estimates scale estimates to be per thousand euro drop in parental DI payments as a result of the reform.[12] Applying this scaling, a parental drop of 1,000 euros results in a 0.9 percentage point lower probability a child will be on DI and 38 fewer cumulative days on DI.

5.2 Other Government Transfer Programs

To arrive at the cost savings to taxpayers from the reduced DI use of children, in Figure 4 we plot an RD graph with the dollar amount of cumulative DI receipt as the dependent variable. Mirroring the drop in total child days on DI at the cutoff, there is a drop of 1,600 euros in cumulative child DI benefits between 1999 and 2014. As reported in Table 3, this is a sizable effect relative to the mean

[12] Note that the IV estimates will have the opposite sign compared to the reduced form, as the first stage estimate is negative.

of 10 thousand euros in DI receipt on average (including zeros). The IV estimates which provide one possible scaling indicate that when a parent's DI benefits fall by a thousand euros, a child's cumulative DI income by 2014 is 1,256 euros lower.

We next look at child participation in other government transfer programs. This is important, because if children are simply shifting from one social assistance program to another, the cost savings to the government from children's reduced DI use will be overstated. Indeed, Borghans et al. (2014) document that while the reform lowered DI participation and benefits for those directly affected, a sizable portion of this loss was replaced by increased participation in other social assistance programs in the short run. In particular, they find a 300 euro increase in other government transfers for each thousand euro lost in DI benefits. Similar program substitution occurs for the directly affected parents in our sample as well.

With this motivation in mind, we pool together 30 miscellaneous benefit programs which are part of the social safety net in the Netherlands, and see if a child's receipt of these other benefits is affected by having a parent subject to the harsher DI rules. The bottom graph in Figure 4 reveals no noticeable change in other benefit receipt at the cutoff. Table 3 present RD regression estimates for these other child benefit outcomes. The point estimate is small and statistically insignificant. The table breaks things down further by separately reporting RD estimates for UI income, general assistance (traditional cash welfare), and the remaining miscellaneous benefit programs. For each type of benefit category, the estimates are small and statistically insignificant.

These results stand in stark contrast to those of their parents, who themselves had substantial substitution to these other programs in the short run (in particular to the UI program). This means that a parent's increased reliance on these other transfer programs, including any accrued knowledge and experience, did not transfer to their children. Any learning and spillover effects are apparently linked to the DI program. The conclusion is that the cost savings for the next generation due to lower DI use is not offset by increased participation in other programs.

5.3 Labor Market Earnings and Taxes Paid

We now turn to labor market earnings and taxes paid by children. The top panel of Figure 5 plots the cumulative earnings of children for the 15 year period from 1999 to 2014. Cumulative earnings includes wage income as well as income from self employment. In this graph, we plot the residuals from a regression of child earnings on child age and age squared. Children's cumulative earnings have a steep own-age profile and child age increases on average with their parent's age as of the reform date. This makes the range of the y axis so wide with raw data that it is difficult to zoom in on the RD jump at the cutoff, which is why we residualize. We note that while child age is positively correlated with the age of their parent, this should not be a problem, as child age appears to be smooth through the parent age cutoff of 45.[13]

The top figure shows a jump at the parental age cutoff in cumulative child earnings. Turning to Table 4, the RD estimate is an increase of a little over 7 thousand euros in earnings for children whose parents were subject to the reform. This is roughly a 2 percent increase in earnings relative to the overall mean. Stated somewhat differently, the IV scaling suggests that for each 1,000 euro drop in parental DI benefits due to the reform, children's cumulative earnings increase by 5,700 euros.

While increased earnings are an important result, what matters for the government's balance sheet is taxes minus transfers. We therefore calculated predicted taxes for children from 1999 to 2014. The bottom graph of Figure 5 plots child cumulative child tax payments versus the running variable of parental age. As we did for earnings, we first regress out a child's age and age squared for this graph. As Table 4 documents, there is a large and statistically significant reduced form effect on taxes: estimated taxes paid rise by two thousand euros, which is a little under two percent of the mean. The IV estimate which provides a scaling is also sizable, but loses statistical significance at the 10 percent confidence level.

5.4 Cumulative Fiscal Effects

To provide a comprehensive picture of the fiscal spillover effects, we now estimate the cumulative change in taxes minus transfers up through 2014. Policy makers should ultimately be concerned with this net effect, since this is what matters for the government's budget. To do this, we create a variable which combines DI and all other government transfer program payments and subtracts this from taxes paid by a child. We find that taxes minus transfers increase by 3,481 euros (s.e. = 1,271) for children of parents who were subject to the stricter DI rules. The scaling provided by our IV estimate implies that for each 1,000 euro loss in parental DI benefits, the government's budget improved by 2,770 euro (s.e = 1,325) per child.

To provide further insight into the fiscal effects over time, Figure 6 plots cumulative DI benefits, cumulative other transfers, and cumulative tax payments over time. There is a small, but statistically significant savings in DI payments almost immediately, and this effect grows progressively larger over time. In contrast, other cumulative transfers are close to zero and insignificant for the entire period.

[13] Estimating an RD regression with child age as the outcome variable, and parent's age as the running variable, yields a small, and statistically insignificant coefficient of -.044 (s.e. = .066) at the cutoff.

Cumulative tax payments, plotted in the upper left graph, start out small and rise little in the first 5 years. This makes sense, as many of the children are still in school and have not yet begun working full time in the early years of our data. But the increase in estimated tax payments rises with time, so that by 2006 the effect becomes statistically significant.

The lower right panel in Figure 6 plots the net effect of taxes minus transfers over time. It mirrors the reduction in DI payments and the rise in tax payments over time, as expected. It is interesting to note that by 2014, increased taxes account for a slightly larger fraction of the net savings to the government's budget compared to the reduction in DI payments. This highlights the limitations of looking at DI in isolation, without considering other possible fiscal spillovers.

6 Spillovers in Education and Other Outcomes

6.1 Educational Investments

So far, we have examined how parents influence their children's participation in DI, other government benefit programs, and earnings from work. These child outcomes mostly occur in the future, after a child has grown up and entered the labor market. Is it possible that children anticipate this lower reliance on DI and increased labor market attachment in the future, and make different investment choices while they are still young?

One way to get at this question is to see if children increase their educational investments in response to having a parent exposed to the harsher DI rules. We collected data on children's educational attainment as of 2014.[14] In Figure 7, the top graph plots child years of education against the running variable of the parent's age as of the reform date. While most children will be done with their formal education by 2014, not all are. Indeed, one can see in the figure that education trends slightly upward in the graph as a function of parental age. Table 5 reports the corresponding RF estimate and standard error for years of education. There is a significant jump at the reform cutoff, with children of reform-exposed parents getting .12 years more education, relative to a mean of 11.5 years. The IV estimate suggests that a one thousand euro loss in parental benefits results in an increase of roughly one tenth of a year of education.

The bottom panel in Figure 7 plots the RD graph with upper secondary school completion (roughly the equivalent of High School) as the outcome variable. There is a significant jump of 2.2 percentage points at the reform cutoff, as documented in Table 5. This is a modestly sized, but

[14]The sample size is somewhat smaller than for the analyses in Section 5, because for earlier cohorts, education is only available for a subsample of observations.

economically significant, effect relative to the overall mean of 78 percent.

Table 5 further reports RD estimates for other levels of schooling. As background, from the ages of roughly 4 or 5 to 12 or 13, children attend lower secondary school. Further education in upper secondary school is split into three tracks, and takes an additional 4 to 6 years depending on whether the student enrolls in a vocational or college preparatory program. Access to college is restricted based on the track a student enrolls in during upper secondary school.

We find no effect of a parent's exposure to the DI reform on their children's completion of lower secondary schooling. This is as expected, since most children are too old to be affected, and most children complete this minimal level of schooling anyway. In contrast, children of reform-exposed parents are not only more likely to complete upper secondary school, but they are also more likely to obtain higher education. This could be in part because admittance to college requires completion of upper secondary school.

These results are intriguing, because they provide the first well-identified and precisely estimated evidence (as far as we know) documenting anticipatory investments by children as a result of parental program participation. These higher levels of education investment have the potential to increase future earnings, lower unemployment spells, and hence increase governmental tax revenue, a theme we return to in Section 8 when we interpret our results.

6.2 Other Outcomes: Crime and Marriage

Before continuing, we briefly explore two other sets of outcomes. We start by looking at whether children's crime is affected by having a parent subject to the stricter DI rules. As Table 6 documents, we find a significant reduction in the chances a child is incarcerated if their parent was exposed to the reform. There is a statistically significant 0.3 percentage point drop in incarceration relative to an overall mean of 1.8 percent, or a 16 percent reduction. However, we find no significant evidence for a decline in arrests, which raises the possibility that the incarceration result is a statistical anomaly. However, another explanation is that parents not on DI are better situated to help their child avoid incarceration, perhaps because they have new contacts at work or increased interactions with other government agency support staff.

Looking at marriage outcomes, we find some evidence that having a parent with reduced DI benefits increases the probability a child will get married. There is a 1.1 percentage point increase in marriage, relative to a base of 46 percent, for children of reform-exposed parents. In contrast, cohabitation which involves a child in the relationship goes the other direction, although it is not statistically significant. Insofar as marriage represents a more stable type of union compared to cohabitation, these are potentially positive spillovers. Finally, the table reports the RD reduced form estimate for spousal income. While the coefficient estimate is positive, it is small and imprecisely estimated.

7 Heterogeneous Effects and Robustness

7.1 Effects by Child Age

To better understand the intergenerational spillovers just documented, in this section we break up the estimated effects by child age. We split children into those who are younger or older than 19 as of November 1996. The reason to focus on child age as of this date is that November 1996 is when the Dutch Parliament decided the 45-50 year old cohort would be grandfathered in under the old DI rules. It is also the approximate time when the re-examinations for the 40-45 year old cohorts began. Hence, November 1996 most accurately captures when children began to be differentially affected by the reform (based on whether their parent was exposed to the stricter DI rules or not).

As a reminder, we limit our sample to children still living at home at the time of the reform announcement. Almost 61,000 children are 18 or younger in our sample as of November 1996, and roughly 55,000 children are 19 or older. Too understand the relative sizes of these samples, one needs to take into account two things. First, since parents are age 40-50 in our estimation sample as of August 1993, they are roughly 43 to 53 by November 1996. Few parents in this age range have young children, since fertility is highest when parents are in their twenties and early thirties. Second, children in the Netherlands commonly live with their parents at older ages.

Table 7 reports RD estimates for our main outcomes split by child age. There are separate regressions for children who are 18 or younger versus 19 or older as of November 1996. Looking at DI participation spillovers in specifications A through C, the effects are all large and statistically significant for the 18 and younger group. The estimated effects for the older age group, while going in the same direction, are smaller and not statistically significant. For example, the reduced form coefficient for the cumulative number of days on DI is twice as large for the younger group. The contrast becomes even starker when comparing the IV estimates which scale the estimates. This is because reform-exposed parents of younger children experience reductions in DI payments which are roughly 25 percent smaller compared to parents of older children (see panels I and J).

Looking at benefits from other programs, excluding DI, we find no effect for either age group.

But when we turn to income from work, we find relatively large and statistically significant effects for the older group. These increased earnings also translate into higher taxes paid. For the younger group, the effects are the same sign, but smaller and not statistically significant. This apparent puzzle, given the opposite pattern found for DI participation by child age, has a simple explanation which we return to after discussing the education outcomes.

Specifications G and H estimate the spillover effects on child education. As background, schooling is compulsory until age 18 in the Netherlands (unless an individual is working). Treated children who are younger receive roughly one-sixth fewer years of education. The IV estimate for the younger group implies an increase of roughly .16 years of schooling for each thousand dollar reduction in parental DI benefits. In contrast, there is no statistically significant spillover in years of education for those age 19 and older. Looking at upper secondary school completion, we again find larger effects for the younger age group. Upper secondary school is usually completed by age 18 or 19, so for the older group, there is less time to affect this schooling margin. We do, however, observe a marginally significant effect for the 19 and older group. This could be due to an effect for children who are in the 6 year educational track, or for children who have previously repeated a grade.[15]

How do all of these results fit together, particularly the pattern of stronger DI effects for the younger group and stronger earnings effects for the older group? First, it is important to recognize that because of their age, the older group has had roughly three more prime-age years to work in the labor market; indeed, mean cumulative earnings for the older group are 50% higher. More importantly, the younger group gets more education, which delays the start of their prime earnings years. Accounting for this education-induced absence from the workforce can more than explain the difference in the earnings effects found for the two age groups.[16] On top of this, there is a return to higher education in the labor market, both in terms of higher wages and less time spent unemployed. This suggests the earnings effect for the younger group should continue to grow over time relative to older group. Notice that education-induced absences from the workforce can also help explain the stronger DI result for younger children, as individuals cannot be enrolled full time in school and concurrently on DI.

In summary, the pattern of results in Table 7 indicates that younger children are generally more

[15]When we restrict the sample to those age 20 or older, we find no statistically significant effect.

[16]To make a comparison, we concentrate on the IV estimates to account for the differences in parental first stages for the two age groups, but a similar point could be made with the RF estimates. Treated children in the younger group receive an extra .158 years of education for each thousand dollars reduction in parental DI, compared to .043 for the older group. A reasonable estimate of earnings in prime age years can be taken from the difference in average cumulative earnings for the two groups (451,868 - 298,005 euros) divided by the average age difference between the groups (2.83 years). Assuming individuals do not work while in school, this implies a loss of 6,252 in earnings for treated children. Adding this to the IV estimate of -3,283 from specification E in Table 7 equals -9,535, which is larger in absolute value compared to the IV estimate of -6,817 for the older age group.

affected by their parent's experience with the DI program. A natural explanation is that younger children are at a more impressionable developmental stage. Not only are younger children's plans for the future are likely more malleable, but a parent's influence is likely larger for minor children. On top of this, younger children have more remaining years living at home to be exposed to their parent's DI experiences. All of these explanations suggest that younger children would be expected to react more strongly to their parent being kicked off of DI or having their benefits cut.

7.2 Robustness and Placebo Tests

Appendix Table A2 reports a variety of specification checks for our main outcomes. For simplicity, the table only reports the reduced form estimates. The first row repeats our baseline estimates for ease of comparison. In Specification B, we allow separate quadratic trends on each side of the cutoff; the estimates are larger, but the standard errors also increase. The next two specifications remove the triangular weights and the control variables from the regression, and yield similar findings to the baseline.

In specifications E and F we narrow the RD estimation window. Our baseline window is 60 months on each side of the cutoff, or five years. As we shrink the window down to 45 or 30 months, the estimates become somewhat larger, but the standard errors increase as well. All estimates remain statistically significant (except for cumulative other transfers, which isn't significant in the baseline specification either). As an alternative set of specifications, we estimate local linear regressions of varying bandwidth in G - I. The results are similar to before, with smaller bandwidths resulting in slightly larger estimates and standard errors. The tax results are estimated imprecisely enough so as not to be statistically significant, but otherwise the findings are similar to the baseline.

In specification J, we run a regression where the sample only includes children not living at home at the time of the reform. This can occur if the child has moved out or because a child does not live with their biological parent (for example, if the parent is divorced). With this sample, we find no significant effect for most of the outcomes. In specification K we cluster the standard errors by parental age and find it makes little difference. The final two specifications exclude non-native Dutch from the estimation sample and children whose parents left DI by 1995. The RD estimates remain similar to the baseline.

To further explore the validity of our estimates, we conduct placebo tests for our main outcomes. To do this, we collected a completely different sample of children: those whose parents were **not** on DI as of 1996. Since these parents are all subject to the new DI examination rules (regardless of their age), they should not be treated differentially. As a result, there should be no discontinuity at the 45 year old age cutoff. Indeed, we find no evidence of a first stage for this sample.

Appendix Table A3 replicates our baseline reduced form specifications for child outcomes, with the only exception being the different sample. There is no evidence of any effect for any of these outcomes, with the point estimates being uniformly small and statistically insignificant. This provides reassurance our results are being driven by the change in DI strictness, and not other policies which differentially affected parents at a similar age cutoff.

8 Discussion

8.1 Comparison to OLS

To help frame the magnitude of our findings, we first compare them to OLS. To construct our OLS estimation sample, we take all parents between the ages of 40 to 50 at the time of the reform who have at least one child still living at home with them. This sample includes parents who were on DI as of the reform date, but additionally includes parents who were not on DI as of the reform date. To make sure differential re-examinations for those under versus over the age 45 cutoff do not drive our OLS estimates, we split this sample into two groups: parents between the ages of 40 and 45, and parents between the ages of 45 and 50. We estimate separate OLS regressions for children whose parents are in the younger versus older age groups.

In Table 8, we estimate the effect of parental DI benefit amounts in 1996 on each of our main child outcomes. We include the same set of control variables as in Table 2. The OLS estimates are most directly comparable to the IV estimates shown previously, as both are measured on the same scale. The OLS estimates for the younger versus older parent samples are generally quite similar, but diverge sharply from the IV estimates.

The first row in Table 8 uses whether the child was ever on DI by 2014 as the outcome variable. The OLS estimate implies an extra 1,000 euros in parental DI payments increases a child's probability of participating in DI by .3 percentage points for both the older and younger parent samples. These OLS estimates compare to the IV estimate of .9 percentage points in Table 2. In other words, the IV estimate is three times larger than the OLS estimate. Looking at cumulative days on DI, cumulative DI income, earnings, taxes, or education, the IV estimate is between 2 and 4 times larger compared to OLS. Interestingly, the OLS estimate for cumulative total benefits, excluding DI payments, is approximately the same magnitude as the OLS estimate for cumulative DI income. In contrast, the

IV estimate of cumulative total benefits are close to zero.

Why are the IV estimates substantially larger? There are several possible explanations. First, the reasons for differential DI participation and payment amounts are likely not the same in the two samples. For example, in the OLS sample, if a child observes a parent is "gaming" the system, they may be inclined to do the same, but if they see their parent is truly disabled, it may have little effect on them unless they experience the same health condition. In contrast, the IV estimate compares parents whose job prospects and health conditions are presumably similar, but whose DI payments change due to differential exposure to the new, stricter DI rules. Being kicked off of DI or having one's payments reduced could represent a large shock to a parent expecting to remain on DI for the long term, and children's views about the ability to rely on government support could change markedly in response.

A second reason is that IV estimates a local average treatment effect (LATE) for compliers, and the intergenerational spillovers could be different for the complier sample compared to the general population. Indeed, the sample means appearing in Table 8 suggest the RD samples are more disadvantaged. For example, average child DI participation and usage is roughly 30 to 40 percent lower in the OLS samples, and years of education is almost one year higher. The relative disadvantage of the RD sample makes sense, as it only includes children of parents who are already on DI.

To better understand who the compliers are in the RD regressions, Appendix Table A4 calculates the average characteristics of compliers.[17] The table then compares these averages to the characteristics of all children in the OLS sample whose parents are between 44.5 and 45.5 years old as of the reform date. The biggest difference is the degree of parental disability. Fifty-seven percent of individuals in the OLS sample are fully disabled, while only 48 percent of compliers are fully disabled. Compliers have also been on DI for a longer time period, with durations which are 10 months longer on average.

8.2 Policy Implications

To be written.

[17] For details on how to calculate the complier averages, see Borghans et al. (2014).

9 Conclusion

Does participation in a social program by a parent influence their child's use of public assistance, employment, and human capital investments? This is a difficult question to answer due to the nonrandom nature of program participation and the likelihood that unobserved factors driving participation are correlated across generations. Yet from a policy perspective, what children learn from their parents about employment versus government assistance could matter for the financial stability of a variety of social insurance and safety net programs.

In this paper we take advantage of a disability insurance (DI) reform in the Netherlands which simultaneously tightened eligibility criteria and reduced the generosity of the program. It is important to recognize the reform affected marginal DI participants who had more work capacity on average. The intergenerational spillover effects for these marginal cases could be quite different compared to cases where a parent has little or no work capacity. Fortunately, the marginal participants we study are the most policy relevant, as they are the ones who presumably would be targeted by most reforms. We find that children whose parents were exposed to the reform are less likely to participate in DI themselves as adults, do not increase their participation in other public assistance programs, increase their earnings, and invest significantly more in their education.

From a policy perspective, our study serves as an important lesson for the evaluation of costs and benefits associated with reforms. Considering just the direct effects on current participants, without accounting for peer effects within families or other networks, would be a mistake. In our setting, the intergenerational spillover effects are substantial, accounting for roughly 20% of the cost savings in the long run; ignoring these spillovers greatly underestimates the cost savings of the Dutch DI reform we study.

References

To be added.

Figure 1: DI Stocks and Inflows as a Percentage of Insured Workers, 1968-2016.

Notes: Data come from the Dutch Employee Insurance Agency (Uitvoeringsinstituut Werknemersverzekeringen), as used in Koning and Lindeboom (2015). Estimates of the number of insured workers are used to calculate receipt and award percentages for 2014 to 2016.

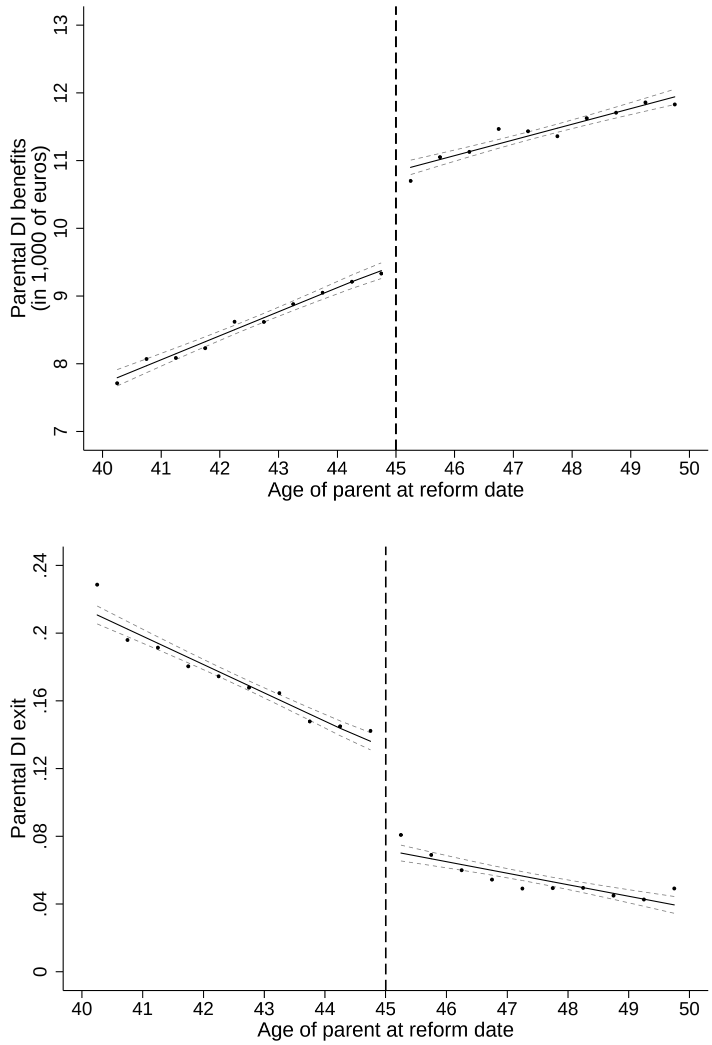

Figure 2: Effects of the Reform on Parents

Notes: Each observation represents average parental DI receipt (top panel) or average parental DI exit (bottom panel) in 6 months age bins, based on the parent's age as of the reform date of August 1993. The dashed vertical lines denote the reform cutoff of age 45. The solid trend lines are based on regressions using unbinned data, with dotted lines indicating pointwise 90 percent confidence intervals.

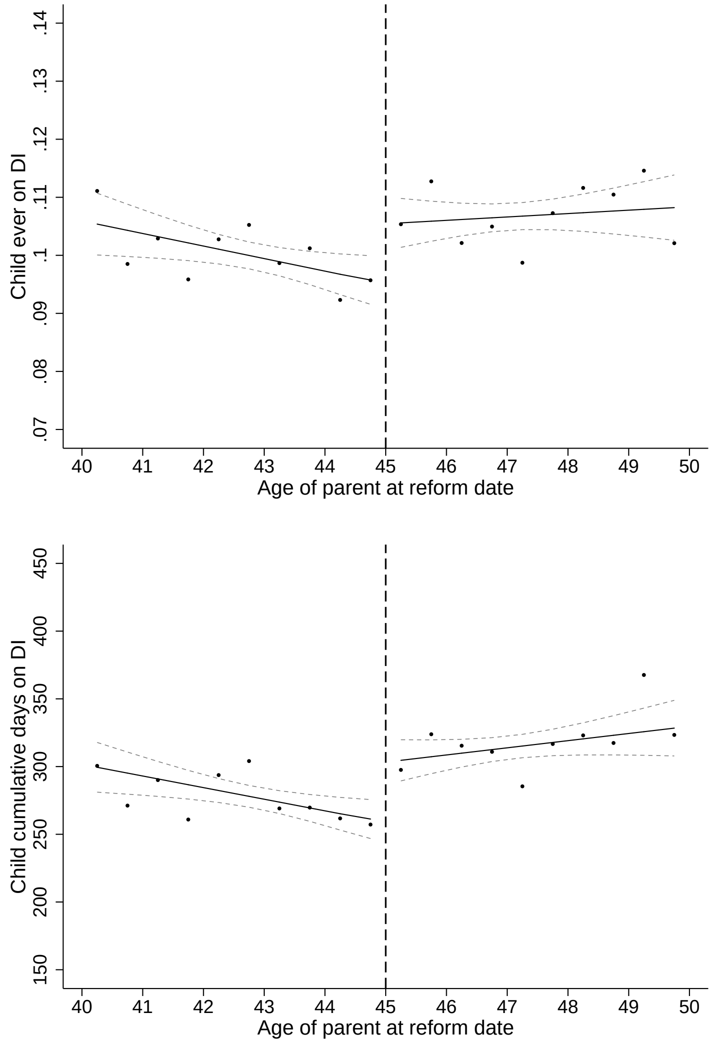

Figure 3: Child DI Participation

Notes: Each observation represents average child DI participation by 2014 (top panel) or average cumulative child days on DI by 2014 (bottom panel) in 6 months age bins, based on the parent's age as of the reform date of August 1993. The dashed vertical lines denote the reform cutoff of age 45. The solid trend lines are based on regressions using unbinned data, with dotted lines indicating pointwise 90 percent confidence intervals.

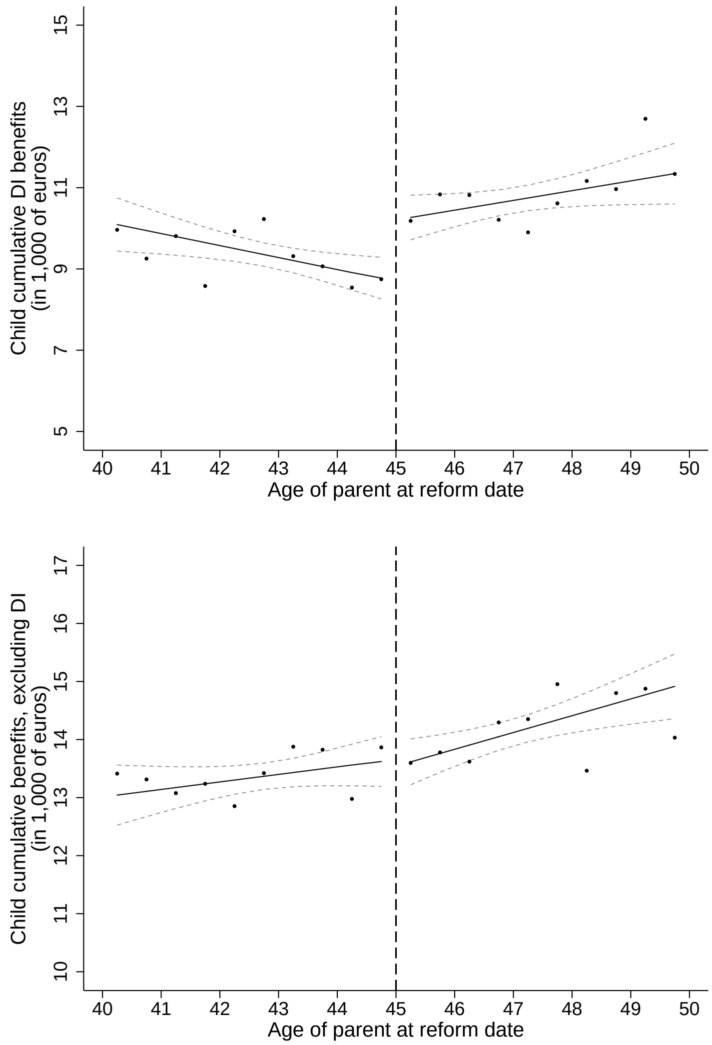

Figure 4: Child DI and Other Benefit Receipt

Notes: See notes to Figure 3.

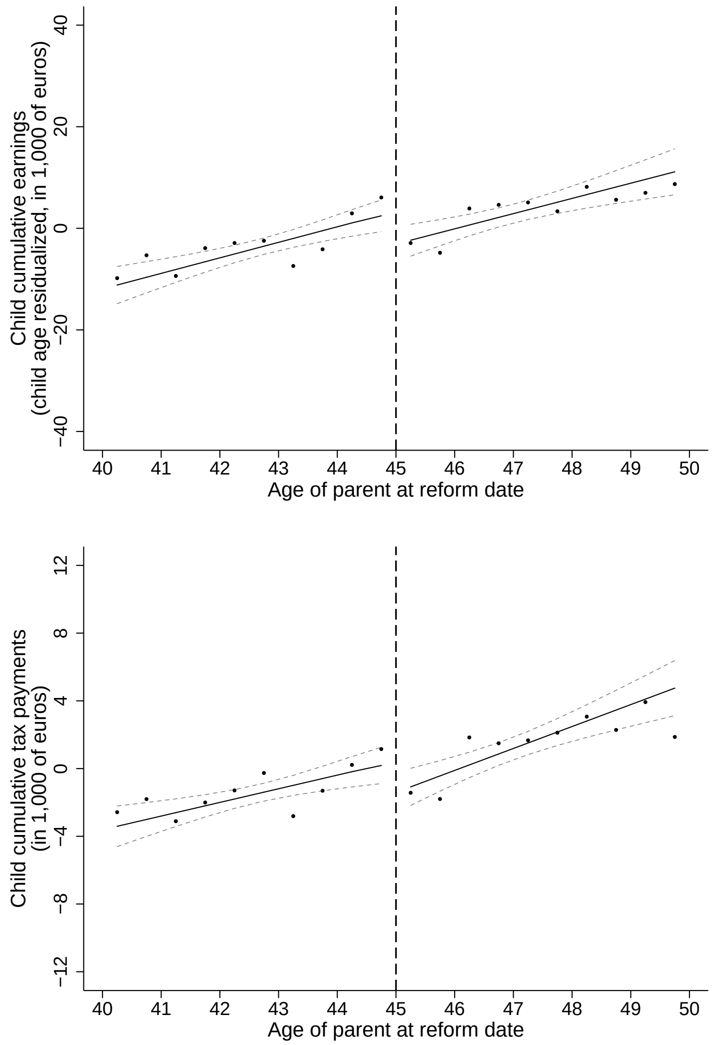

Figure 5: Residualized Child Earnings and Taxes

Notes: See notes to Figure 3. In these graphs, we first regress out child age and age squared and then average the residuals in 6 month age bins.

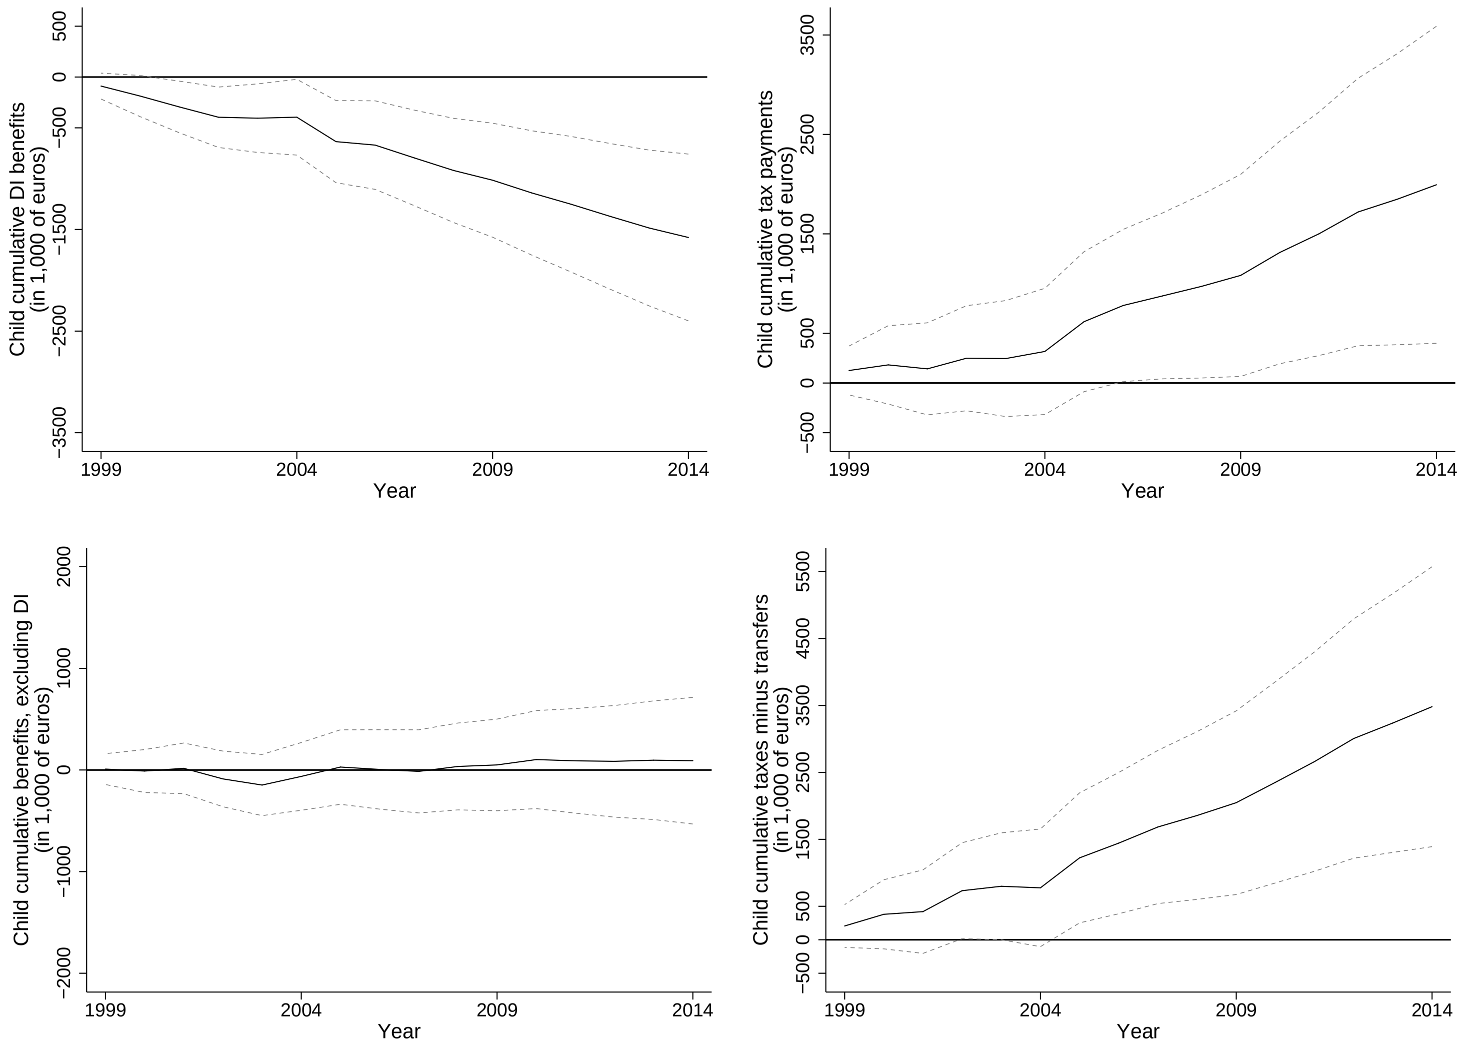

Figure 6: Cumulative Fiscal Effects Over Time

Notes: Each graph plots year-by-year RD estimates of cumulative effects up to a given point in time, using the baseline specification of Table 2.

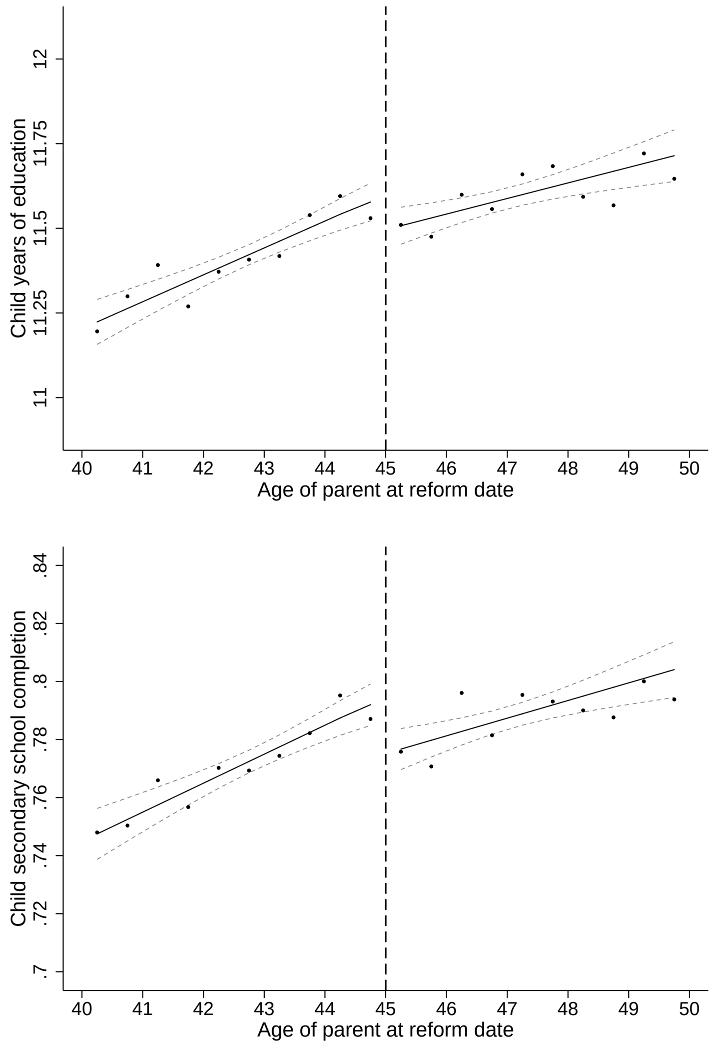

Figure 7: Child Educational Attainment

Notes: See notes to Figure 3.

Table 1: First Stage RD Estimates of the Reform on Parental DI

Dependent variable	Mean	First Stage
A. Parental DI benefits	10.063	-1.300**
		(.115)
B. Parental exit from DI	0.114	0.054**
		(0.005)
Observations		116,356

Notes: Parental DI benefits measure payments received in 1999, in thousands of euros. The sample is parents age 40-50 on August 1, 1993, receiving DI benefits on August 1, 1993, with children living at home on August 1, 1993, and who where still on DI as of January 1, 1996. Parents with multiple children appear more than once in the sample, while children with two parents on DI are dropped from the sample. All coefficients are estimated using an RD model with separate linear trends on each side of the cutoff and triangular weights. The running variable in the RD is the age of the parent as of August 1, 1993 and the cutoff is age 45 as of August 1, 1993. Control variables for the parent are measured as of 1996 and include age, birth month dummies, a gender dummy, a cubic in pre-disability earnings, six dummies for degree of disability, a cubic in DI duration, national origin dummies; control variables for the child are measured as of 1996 and include age, birth month dummies, and a gender dummy. 1996 is before the passage of the law exempting the age 45-50 cohort and before the DI re-examinations for the age 40-45 cohort take place. Standard errors in parentheses, clustered at the parent level.

$^{**}p<0.05$, $^{*}p<0.10$

Table 2: Child DI Participation

Child outcome in 2014	Mean	RF	IV
A. Ever on DI	.104	-.011**	.009**
		(.004)	(.004)
B. Cumulative days on DI	298	-47.2**	37.5**
		(13.9)	(14.6)
Observations	116,356		

Notes: See notes to Table 1 for the sample definition, the RD estimator, and the included control variables. The independent variables measure whether a child ever participated in DI between 1996 and 2014, the cumulative number of days on DI between 1996 and 2014. Standard errors in parentheses, clustered at the parent level.

$^{**}p<0.05$, $^{*}p<0.10$

Table 3: Child Benefits from DI and Other Government Programs

Child outcome in 2014 (in 1,000 euros)	Mean	RF	IV
A. DI benefits			
A1. Cumulative DI income	10.107	-1.579**	1.256**
		(.499)	(.522)
B. Other benefits			
B1. Cumulative UI income	5.639	-.067	.053
		(.162)	(.167)
B2. Cumulative general assistance income (traditional cash welfare)	4.432	.092	-.073
		(.266)	(.272)
B3. Cumulative misc. benefit income (all other government safety net programs)	3.675	.067	-.053
		(.145)	(.148)
B4. Cumulative total benefits, excluding DI (B1+B2+B3)	13.746	.091	-.073
		(.379)	(.388)
Observations	116,356		

Notes: See notes to Table 1 for the sample definition, the RD estimator, and the included control variables. All of the independent variables measure cumulative income amounts between 1996 and 2014 for the child, measured in 1,000 euros. Misc. benefit income does not include subsidies, such as those for child care. Standard errors in parentheses, clustered at the parent level.
$^{**}p<0.05$, $^{*}p<0.10$

Table 4: Child Earnings and Taxes

Child outcome in 2014 (in 1,000 euros)	Mean	RF	IV
A. Cumulative income from work	371.282	7.172**	-5.706*
		(2.836)	(2.951)
B. Cumulative estimated taxes	109.565	1.994**	-1.587
		(.969)	(1.008)
Observations	116,356		

Notes: See notes to Table 1 for the sample definition, the RD estimator, and the included control variables. All of the independent variables measure cumulative income amounts between 1996 and 2014 for the child, measured in 1,000 euros. Cumulative income from work includes both wage income and self-employment income. Standard errors in parentheses, clustered at the parent level.
$^{**}p<0.05$, $^{*}p<0.10$

Table 5: Child Educational Investments

Child outcome in 2014	Mean	RF	IV
A. Years of education	11.49	.117**	-.096**
		(.050)	(.050)
B. Lower secondary school or more	.95	-.001	.0009
		(.003)	(.003)
C. Upper secondary school or more	.78	.022**	-.018**
		(.007)	(.007)
D. Bachelor degree or more	.33	.017**	.-014*
		(.008)	(.008)
E. Master degree or more	.10	.009*	-.008
		(.005)	(.005)
F. Advanced degree or more	.01	-.001	.001
		(.001)	(.001)
Observations	79,924		

Notes: See notes to Table 1 for the sample definition, the RD estimator, and the included control variables. Education is measured as of 2014. Upper secondary school or more includes vocational school. The sample size in this table is smaller, as education data is complete for younger cohorts, but incomplete for older cohorts. Standard errors in parentheses, clustered at the parent level.
***p<0.05, *p<0.10*

Table 6: Child Crime and Marriage Outcomes

Child outcome in 2014	Mean	RF	IV
A. Ever arrested	.135	-.002	.002
		(.004)	(.004)
B. Ever incarcerated	.018	-.0030**	.0024
		(.0017)	(.0017)
C. Ever married	.458	.0107**	-.0084
		(.0062)	(.0063)
D. Ever cohabiting (with a child)	.300	-.0049	.0038
		(.0058)	(.0059)
E. Spousal income (including 0's, in 1,000 euros)	112,657	.887	-.697
		(2.406)	(2.445)
Observations	123,186		

Notes: See notes to Table 1 for the sample definition, the RD estimator, and the included control variables. Arrest and incarceration data come from two different datasets. Standard errors in parentheses, clustered at the parent level.
$^{**}p<0.05$, $^{*}p<0.10$

Table 7: Effects by Age of Child

	Child age: 18 and younger			Child age: 19 and older		
Child outcome in 2014	Mean	RF	IV	Mean	RF	IV
A. Ever on DI	.093	-.017**	.016**	.115	-.007	.005
		(.005)	(.006)		(.006)	(.005)
B. Cumulative days on DI	265	-65.9**	60.9**	334	-32.5	22.3
		(18.7)	(22.3)		(21.1)	(16.6)
C. Cumulative DI income (in 1,000 euro)	8.717	-2.283**	2.110**	11.635	-1.121	.769
		(.634)	(.757)		(.785)	(.621)
D. Cum. total benefits, excl. DI (in 1,000 euro)	11.755	-.369	.341	15.936	.388	-.266
		(.457)	(.520)		(.613)	(.478)
E. Cumulative income from work (in 1,000 euro)	298.005	3.552	-3.283	451.868	9.944**	-6.817*
		(3.048)	(3.485)		(4.638)	(3.672)
F. Cumulative estimated taxes (in 1,000 euro)	82.726	391	-361	139.082	3.271**	-2.242*
		(.970)	(1.099)		(1.661)	(1.320)
G. Years of education	11.58	.165**	-.158**	11.36	.063	-.043
		(.065)	(.078)		(.072)	(.053)
H. Upper secondary school or more	.777	.025**	-.024**	.781	.018*	-.012
		(.009)	(.010)		(.010)	(.008)
I. First Stage: Parental DI benefits (in 1,000 euro, for A-F)		-1.082			-1.459	
		(.129)			(.112)	
J. First Stage: Parental DI benefits (in 1,000 euro, for G, H)		-1.043			-1.475	
		(.143)			(.141)	
Observations (A-F)	60,942			55,414		
Observations (G, H)	48,398			31,526		

Notes: See notes to Tables 1-5. Child age is measured as of November 1996.
***p<0.05, *p<0.10*

Table 8: OLS Estimates

	Indep. var.: Parental DI payments in 1996 (in 1,000 euros)			
	Parent age: 40-45		Parent age: 45-50	
Child outcome in 2014	Mean	OLS	Mean	OLS
A. Ever on DI	.060	.003**	.070	.003**
		(.000)		(.000)
B. Cumulative days on DI	164	8.7**	203	8.6**
		(.3)		(.3)
C. Cumulative DI income (in 1,000 euro)	5.495	.298**	7.039	.296**
		(.009)		(.010)
D. Cumulative total benefits, excluding DI (in 1,000 euro)	9.160	.247**	10.707	.246**
		(.007)		(.008)
E. Cumulative income from work (in 1,000 euro)	347.254	-1.842**	442.195	-2.552**
		(.062)		(.080)
F. Cumulative estimated taxes (in 1,000 euro)	99.245	-.489**	134.848	-.782**
		(.023)		(.032)
G. Years of education	12.39	-.053**	12.47	-.051**
		(.001)		(.001)
H. Upper secondary school or more	.85	-.005**	.87	-.005**
		(.000)		(.000)
Observations (A-F)	498,378		421,731	
Observations (G, H)	387,264		287,799	

Notes: Sample includes children still living at home for all parents between the ages of 40 to 50, regardless of whether the parent was on DI as of the reform date. The sample is split into two parental age groups (40-45 and 45-50 as of the reform date) to ensure that differential re-examinations for those parents under versus over the age 45 cutoff do not drive the OLS estimates. Control variables include those used in Table 1. Standard errors in parentheses, clustered at the parent level.
***p<0.05, *p<0.10*

Appendix Figures and Tables

"Intergenerational Spillovers in Disability Insurance"

Gordon B. Dahl and Anne C. Gielen

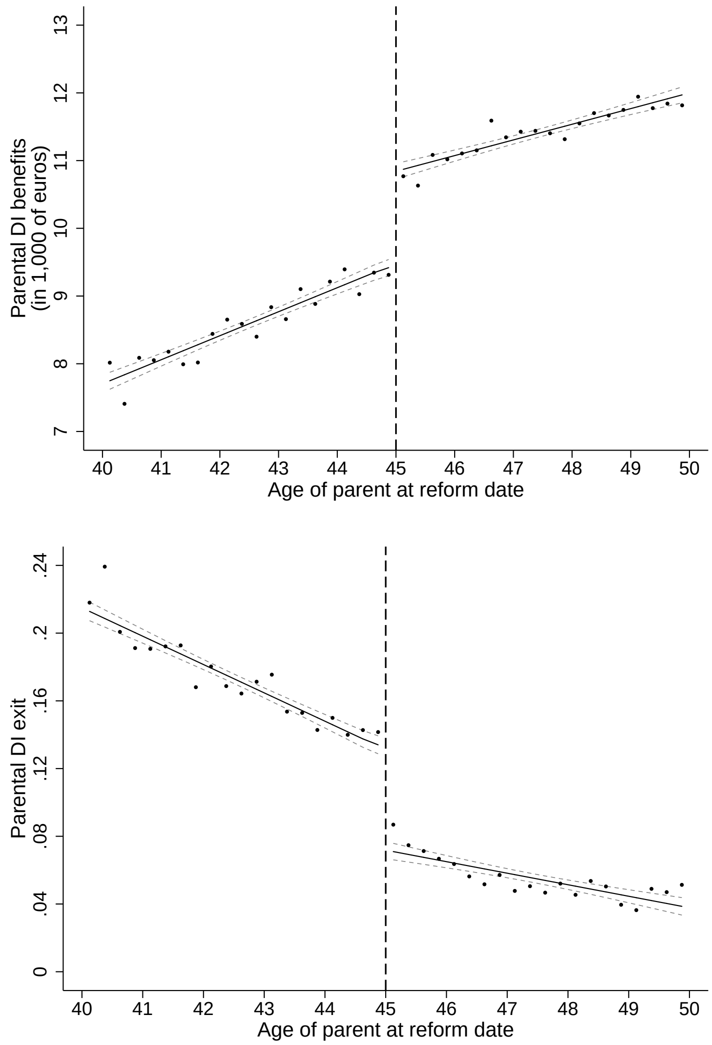

Appendix Figure A1: Effects of the Reform on Parents

Note: Graphs mirror those in Figure 2, but with 3 month age bins.

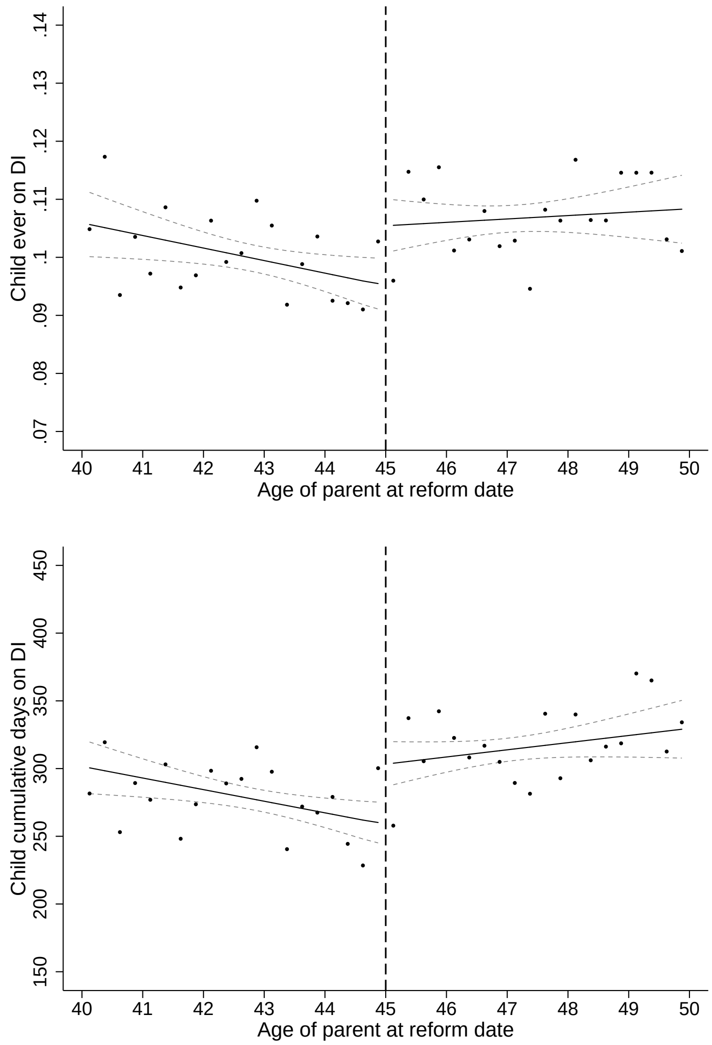

Appendix Figure A2: Child DI Participation

Note: Graphs mirror those in Figure 3, but with 3 month age bins.

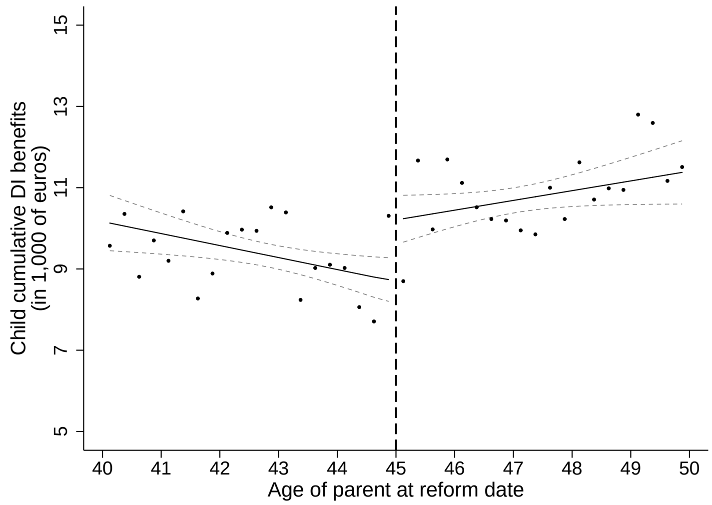

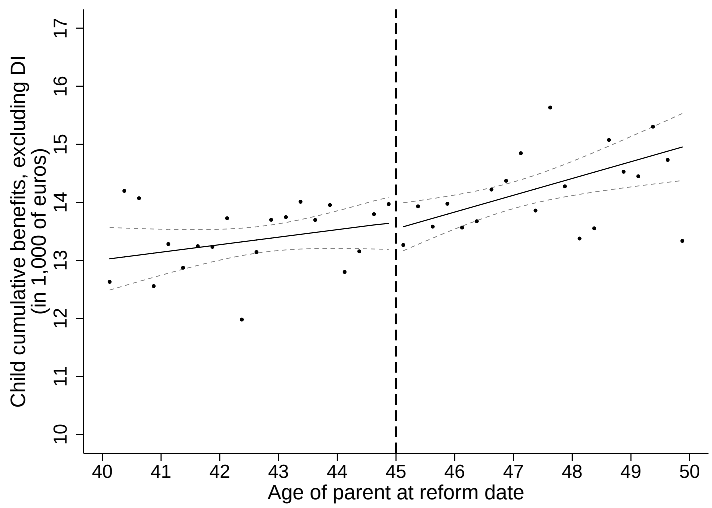

Appendix Figure A3: Child DI and Other Benefit Receipt

Note: Graphs mirror those in Figure 4, but with 3 month age bins.

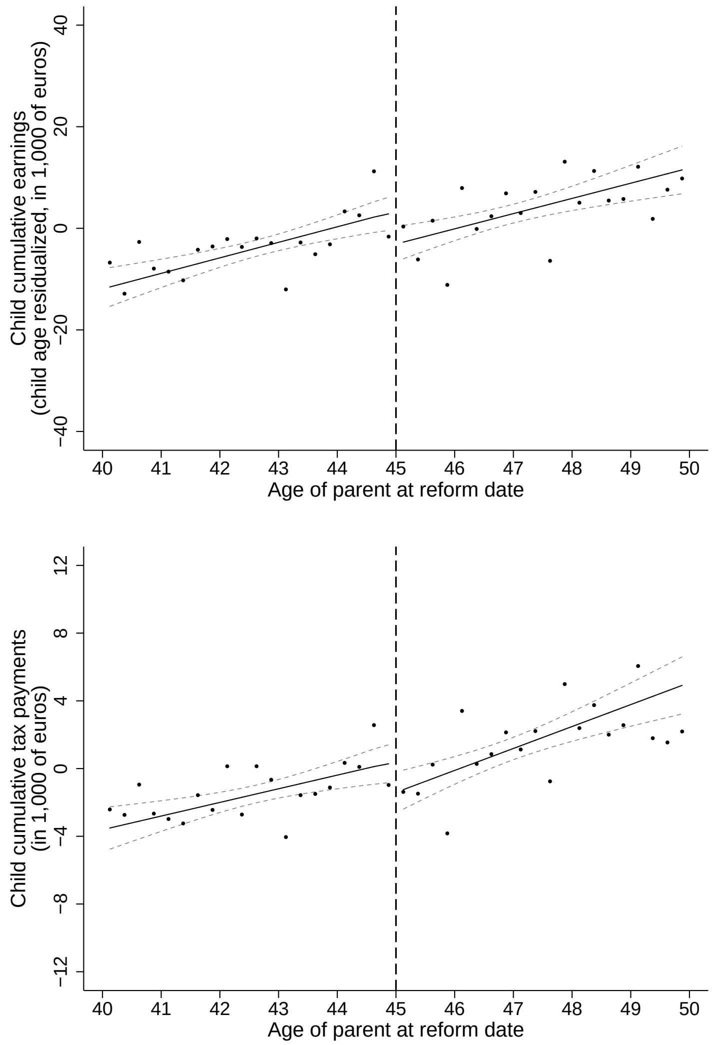

Appendix Figure A4: Residualized Child Earnings and Taxes

Note: Graphs mirror those in Figure 5, but with 3 month age bins.

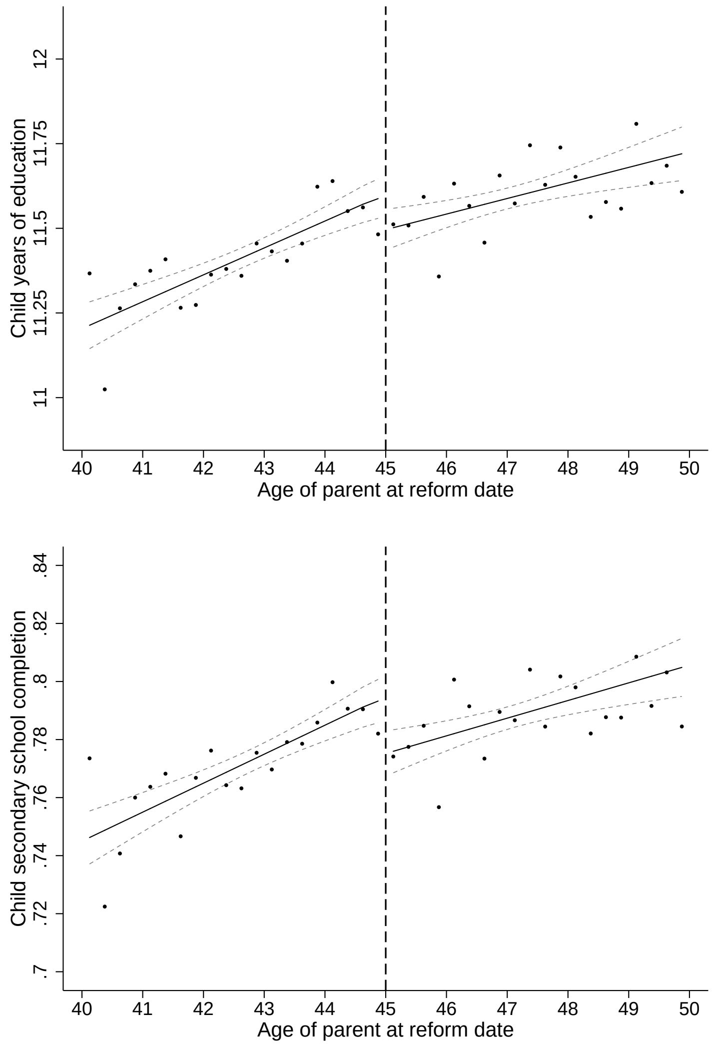

Appendix Figure A5: Child Educational Attainment

Note: Graphs mirror those in Figure 7, but with 3 month age bins.

Appendix Table A1: Summary Statistics

	Overall	Parent age: 40-45	Parent age: $45 \leq$ and $<$ 50
A. Parents			
Female	0.27	0.29	0.26
Married (Jan 1996)	0.87	0.87	0.87
Age (Aug 1993)	45.17	42.58	47.36
Duration DI (Aug 1993)	88.38	85.20	91.08
Degree DI			
2	0.10	0.14	0.07
3	0.12	0.14	0.10
4	0.08	0.09	0.08
5	0.07	0.06	0.08
6	0.02	0.02	0.03
7	0.02	0.02	0.03
8 (Full)	0.58	0.53	0.63
Pre-DI earnings	6529.06	6249.30	6766.46
Native Dutch	0.91	0.91	0.91
>5 years on DI	0.56	0.54	0.57
Number of kids in HH	1.71	1.87	1.58
Parent observations	70,319	32,279	38,040
B. Children			
Female	0.44	0.46	0.41
Age (Aug 1993)	15.16	13.42	16.83
Child observations	116,356	57,028	59,328

Notes: The sample is parents age 40-50 on August 1, 1993, receiving DI benefits on August 1, 1993, with children living at home on August 1, 1993, and who where still on DI as of January 1, 1996. Parents with multiple children appear more than once in the sample, while children with two parents on DI are dropped from the sample.
***p<0.05, *p<0.10*

Appendix Table A2: Robustness Tests for Main Child Outcomes (Reduced Form Models)

Specification	Ever on DI	Cum. days on DI	Cum. DI Income	Cum. other transfers	Cum. earnings	Cum. taxes	Years educ.	Upper Second. School
A. Baseline	-.011**	-47.156**	-1.579**	.091	7.172**	1.944**	.117**	.022**
	(.004)	(13.921)	(.499)	(.379)	(2.836)	(.969)	(.050)	(.007)
B. Quadratic trends	-.016**	-57.876**	-2.075**	.347	10.492**	3.128**	.260**	.037**
	(.006)	(20.334)	(.729)	(.562)	(4.165)	(1.415)	(.126)	(.017)
C. No triangular weights	-.010**	-36.906**	-1.175**	.020	5.269**	1.309	.099**	.018**
	(.004)	(13.054)	(.469)	(.354)	(2.642)	(.908)	(.046)	(.006)
D. No control variables	-.010**	-44.400**	-1.515**	.098	5.877**	1.548	.100**	.019**
	(.004)	(13.822)	(.497)	(.385)	(3.242)	(1.117)	(.052)	(.007)
E. 45 month window	-.012**	-52.185**	-1.776**	.128	8.019**	2.245**	.120**	.022**
	(.004)	(15.094)	(.541)	(.411)	(3.084)	(1.051)	(.054)	(.007)
F. 30 month window	-.018**	-69.736**	-2.439**	.264	11.085**	2.944**	.142**	.026**
	(.005)	(18.142)	(.651)	(.495)	(3.684)	(1.254)	(.065)	(.009)
G. Local linear regression	-.010**	-37.323**	-1.227**	-.019	3.923	.676	.076	.015**
bandwidth 60 months	(.004)	(12.737)	(.459)	(.345)	(3.040)	(1.029)	(.048)	(.006)
H. Local linear regression	-.010**	-46.517**	-1.589	.954	6.116*	1.641	.091*	.018**
bandwidth 45 months	(.004)	(14.526)	(.543)	(.412)	(3.424)	(1.196)	(.055)	(.007)
I. Local linear regression	-.016**	-67.382**	-2.338**	.308	8.993**	2.298	.147**	.025**
bandwidth 30 months	(.005)	(17.858)	(.629)	(.501)	(4.161)	(1.416)	(.067)	(.009)
J. Sample of children not	-.007	-17.811	-.323	-.199	6.039	1.666	.121**	.013
living at home	(.006)	(20.977)	(.812)	(.712)	(4.179)	(1.470)	(.066)	(.009)
K. Cluster s.e.'s	-.011**	-47.156**	-1.579**	.091	7.172**	1.994**	.117**	.022**
by parental age	(.003)	(13.002)	(.512)	(.332)	(2.583)	(.793)	(.043)	(.005)
L. Excluding non-natives	-.011**	-42.127**	-1.374**	.034	8.136**	2.233**	.124**	.023**
	(.004)	(14.629)	(.521)	(.387)	(2.959)	(1.019)	(.052)	(.007)
M. Excluding children whose	-.012**	-49.768**	-1.694**	.014	7.172**	1.837**	.126**	.022**
parents left DI in 1995	(.004)	(14.144)	(.506)	(.384)	(2.866)	(.977)	(.050)	(.007)

Notes: See notes to Tables 1-5. Standard errors in parentheses, clustered at the parent level.
***p<0.05, *p<0.10*

Appendix Table A3: Placebo Tests: RD Estimates for Parents Not on DI in 1996

Child outcome in 2014	Mean	RF
A. Ever on DI	.057	-.0002
		(.0009)
B. Cumulative days on DI	158	-2.130
		(3.232)
C. Cumulative DI income	5.346	-.150
(in 1,000 euro)		(.117)
D. Cumulative total benefits, excluding DI	8.761	.042
(in 1,000 euro)		(.092)
E. Cumulative income from work	378.393	-.060
(in 1,000 euro)		(.940)
F. Cumulative estimated taxes	110.906	-.185
(in 1,000 euro)		(.361)
G. Years of education	12.56	.010
		(.005)
H. Upper secondary school or more	0.87	.0000
		(.0016)
Observations (A-F)	1,286,355	
Observations (G, H)	971,599	

*Notes: The placebo sample is comprised of children whose parents were **not** on DI as of 1996. Since these parents are all subject to the new DI examination rules (regardless of their age), there should be no discontinuity at the cutoff in any of the child outcomes. See notes to Tables 1-5 for details on the RD estimator, the included control variables, and the child outcome variables. Standard errors in parentheses, clustered at the parent level.*

$^{**}p<0.05$, $^{*}p<0.10$

Appendix Table A4: Characteristics of Compliers

	Compliers	Ave at age 45	Difference
A. Parents			
Female	.208**	.226**	-.018
	(.016)	(.003)	(.016)
Married (Jan 1996)	.927**	.896**	.032**
	(.013)	(.002)	(.012)
Months on DI (Aug 1993)	98.697	88.943**	9.754**
	(2.804)	(.432)	(2.760)
Degree DI			
2	.100**	.086**	.014
	(.014)	(.002)	(.014)
3	.077**	.118**	-.040**
	(.015)	(.002)	(.015)
4	.166**	.094**	.072**
	(.014)	(.002)	(.013)
5	.099**	.076**	.023*
	(.012)	(.002)	(.012)
6	.035**	.026**	.010
	(.008)	(.001)	(.008)
7	.041**	.027**	.015**
	(.007)	(.001)	(.007)
8	.481**	.574**	-.093**
	(.020)	(.003)	(.020)
Pre-DI earnings	6,586.082**	6,723.887**	-137.806
	(151.404)	(22.464)	(149.599)
Native Dutch	.910**	.916**	-.006
	(.011)	(.002)	(.011)
Number of kids in HH	2.051**	2.078**	-.027
	(.042)	(.006)	(.041)
B. Children			
Female	.450**	.442**	.007
	(.021)	(.003)	(.021)
Age (Aug 1993)	15.269**	15.194**	.075
	(.183)	(.028)	(.179)

***p<0.05, *p<0.10*