

IZA DP No. 8265

Family Welfare Cultures

Gordon B. Dahl
Andreas Ravndal Kostøl
Magne Mogstad

June 2014

Family Welfare Cultures

Gordon B. Dahl

*University of California San Diego,
Statistics Norway, NBER and IZA*

Andreas Ravndal Kostøl

Statistics Norway and University of Bergen

Magne Mogstad

*University of Chicago,
Statistics Norway, NBER and IZA*

Discussion Paper No. 8265

June 2014

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Family Welfare Cultures^{*}

Strong intergenerational correlations in various types of welfare use have fueled a long-standing debate over whether welfare receipt in one generation causes welfare participation in the next generation. Some claim a causal relationship in welfare receipt across generations has created a culture in which welfare use reinforces itself through the family. Others argue the determinants of poverty or poor health are correlated across generations, so that children's welfare participation is associated with, but not caused by, parental welfare receipt. However, there is little empirical evidence to sort out these claims. In this paper, we investigate the existence and importance of family welfare cultures in the context of Norway's disability insurance (DI) system. To overcome the challenge of correlated unobservables across generations, we take advantage of random assignment of judges to DI applicants whose cases are initially denied. Some appeal judges are systematically more lenient, which leads to random variation in the probability a parent will be allowed DI. Using this exogenous variation, we find strong evidence that welfare receipt in one generation causes welfare participation in the next generation: when a parent is allowed DI, their adult child's participation over the next five years increases by 6 percentage points. This effect grows over time, rising to 12 percentage points after ten years. While these findings are specific to our setting, they serve to highlight that welfare reforms can have long-lasting effects on program participation, since any original effect on the current generation could be reinforced by changing the participation behavior of their children as well. The detailed nature of our data allows us to compare the intergenerational transmission with spillover effects in other networks and to explore mechanisms. Our findings point to a special link between parents and their children, with little impact due to close neighbors' DI participation. We find suggestive evidence that what may change as a result of a parent being allowed DI is their children's beliefs about the efficacy of trying to get on to the DI program or their attitudes about DI participation and its stigma.

JEL Classification: I38, J62, H53

Keywords: intergenerational welfare transmission, welfare cultures, disability insurance

Corresponding author:

Gordon B. Dahl
Department of Economics
University of California, San Diego
9500 Gilman Drive #0508
La Jolla, CA 92093-0508
USA
E-mail: gdahl@ucsd.edu

^{*} We thank four anonymous referees, the editor and seminar participants at several universities and conferences for valuable feedback and suggestions. We are grateful to Knut Brufoss, Espen Vihle and Runar Narvland for their help in accessing the data and in understanding the institutional details. The project received generous financial support from the Norwegian Research Council.

1 Introduction

A large body of evidence demonstrates strong intergenerational correlations in the use of various types of welfare programs, including social insurance and safety net programs. These correlations have fueled a long standing debate over whether welfare receipt in one generation causes welfare participation in the next generation. Some policymakers and researchers have argued that a causal relationship exists, creating a culture in which welfare use reinforces itself through the family. Others argue the determinants of poverty or poor health are correlated across generations in ways which have nothing to do with a welfare culture, but which nonetheless translate into similar participation rates within families. This explanation says that while a child’s use of welfare may be correlated with a parent’s, it is not caused by the parent’s welfare receipt.

Estimating whether welfare receipt in one generation causes welfare participation in the next generation has proven difficult given the likelihood of correlated unobservables across generations.¹ On top of this, it is difficult to access large datasets on welfare use which link family members together across generations. These empirical challenges have meant that existing research has largely focused on intergenerational correlations in various types of welfare use. Black and Devereux (2011), in their Handbook of Labor Economics chapter, summarize the state of the literature well: “while the intergenerational correlations in welfare receipt are clear, there is much less evidence that a causal relationship exists.”

In this paper, we investigate the existence and importance of family welfare cultures, where the receipt of a welfare program by one generation causes increased participation in the next generation. We exploit a policy which randomizes the probability that parents receive welfare in combination with a unique source of population panel data. We investigate the causal relationship in welfare receipt across generations in the context of Norway’s disability insurance (DI) system. Our focus on DI receipt is highly policy relevant, as it is now one of the largest transfer programs in most industrialized countries. In the U.S., for example, outlays for DI exceed those for food stamps, traditional cash welfare, or the EITC.² For families without small children, DI is often the only cash benefit available after unemployment benefits run out and it has therefore become an increasingly important component of the social safety net. Over the past 50 years,

¹Researchers have documented strong intergenerational patterns for a variety of socioeconomic variables (see e.g. Black and Devereux, 2011; Lee and Solon, 2009; Mazumder, 2005; Oreopoulos, Page, and Stevens, 2006), highlighting the difficulty of separating out correlations within families from causal effects. Bjorklund, Lindahl, and Plug (2006) show that both pre- and postbirth factors contribute substantially to the intergenerational transmission of socioeconomic variables. Levine and Zimmerman (1996) show a large portion of the observed correlation in AFDC participation can be explained by intergenerational correlations in income and other family characteristics. Pepper (2000) illustrates the difficulty in drawing causal inferences about intergenerational welfare transmission from observational data.

²In 2011 the U.S. paid out \$129 billion to 10.6 million disabled workers and their families, with an additional \$33 billion worth of disability benefits from the SSI program for poor Americans and \$90 billion in Medicaid for disabled workers (OASDI Trustees Report, 2012). By way of comparison, in the U.S. in 2011 the cash assistance portion of TANF paid out \$10 billion to 4.6 million participants, SNAP (food stamps) paid out \$80 billion to 46.5 million participants and the EITC paid out \$62 billion to 27 million working families. In 2009, DI payments constituted 1.8% of GDP in the U.S. and 2.3% of GDP across the European OECD-countries (OECD, 2010).

DI rolls have steadily risen from less than 1% to over 5% of the adult population in the U.S., from 1% to 7% in the U.K, and from 2% to almost 10% in Norway. Many have argued these increases are fiscally unsustainable, especially as current DI recipients are younger and have longer life expectancies on average compared to previous cohorts of recipients (e.g., Autor and Duggan, 2006; Burkhauser and Daly, 2012).

The key to our research design is that the DI system in Norway randomly assigns judges to DI applicants whose cases are initially denied. Some appeal judges are systematically more lenient, which leads to random variation in the probability an individual will be allowed DI. We utilize this exogenous variation to examine whether parents being allowed DI during the appeal process affects the probability their adult children subsequently apply for and are awarded DI. Our approach takes advantage of the fact that appeal judges are randomly assigned; as a result, the leniency of parents' judges is unrelated to any other intergenerational factors, such as poverty or health, which might influence the DI participation of their children. A similar identification approach based on the quasi-random assignment of judges (or examiners) has been used in other contexts, such as to study the labor supply effects of DI receipt (Maestas, Mullen, and Strand, 2013; French and Song, 2013), the impacts of incarceration (Kling, 2006; Aizer and Doyle, 2013), the consequences of foster care (Doyle, 2007, 2008), and the effects of consumer bankruptcy protection (Dobbie and Song, 2013).

As our measure of judge leniency, we use the average allowance rate in all the other cases a judge has handled. This leniency measure is highly predictive of the judge's decision in the current case, but as we document, uncorrelated with observable case characteristics. Using this random variation as an instrument, we find that when a parent is allowed DI because of a lenient judge, their adult child's participation rate increases by 6 percentage points over the next five years. This intergenerational welfare transmission amplifies over time; the effect of parental DI allowance on their adult child's participation rate reaches 12 percentage points ten years after the judge's decision. By comparison, we calculate that around 3 percent of these children would have been on DI if their parents had been denied DI. Consistent with this increase in adult children's welfare use, we find that parental DI receipt decreases the probability that a child will work or pursue higher education. To assess the internal validity of our research design, we perform a number of robustness checks, all of which suggest the identifying assumptions of independence, exclusion and monotonicity hold.

As in Bertrand, Luttmer, and Mullainathan (2000), we think of the spillover effects in welfare receipt within families or other social networks as measures of welfare culture, with the understanding that culture may operate through information, beliefs or norms. Our rich data allows us to take several steps to explore the breadth and nature of such welfare cultures. First, we go beyond the transmission of DI receipt across generations and use our research design to examine spillovers in other social networks. Our findings point to a special link between parents and their children, with little, if any, impact of close neighbors' DI receipt. By

comparison, we do not have enough precision to draw firm conclusions about spillovers in DI receipt across siblings or spouses. Second, we examine how the intergenerational transmission of DI receipt depends on the type of parent-child link. Our findings suggest that parents' influence on childrens' decisions to apply for and take up DI is not specific to the living arrangement or age of the child. Third, we explore how intergenerational welfare transmission could operate in our context. By construction, it is unlikely to be information about how to apply for the program or appeal an unfavorable decision, since all parents in our dataset go through these two processes. We find suggestive evidence that what may change as a result of a parent being allowed DI is their children's beliefs about the efficacy of trying to get on to the DI program or their attitudes about DI participation and its stigma. Part of this evidence comes from an analysis which shows that children whose parents got a lenient judge are not only more likely to apply for DI, but are also more inclined to report the same type of medical disorder.

Our paper complements a growing literature on the causes and consequences of the growth in DI rolls (for a review, see Autor and Duggan, 2006; Autor, 2011). To date, research has largely focused on estimating the work capacity and labor supply elasticity of DI recipients.³ Yet despite a recent surge in research on this topic, less is known about what causes individuals to apply for DI, why disability rolls have risen so dramatically, and how the receipt of DI affects individuals on margins other than labor force participation.⁴ Our study provides some of the first causal evidence on what influences DI applications and what the effects of DI receipt by parents are for their children. The magnitude of our estimates suggest that intergenerational transmission could play a role in explaining the dramatic rise in DI rolls over the past few decades.

Two studies using U.S. data and a similar research design have looked at how DI receipt affects labor supply. Maestas, Mullen, and Strand (2013) use variation in the leniency of initial examiners in the U.S. and find that DI receipt substantially reduces earnings and employment of applicants. Exploiting the leniency of appeal judges in the U.S., French and Song (2013) find comparable labor supply effects of DI receipt among appellants. What makes our study unique is the ability to link the judicial decisions to a wide range of variables for both parents and their children. This allows us to provide novel evidence on intergenerational welfare transmission in a setting where we can credibly address concerns about correlated unobservables across generations.

At the same time, it is important to emphasize the local nature of our results. Our IV estimates are a local average treatment effect (LATE) for children whose parents could have received a different allowance

³See e.g. Autor and Duggan (2003); Borghans, Gielen, and Luttmer (2013); Bound (1989); Campolieti and Riddell (2012); French and Song (2013); Gruber (2000); Kostol and Mogstad (2014); Maestas, Mullen, and Strand (2013); Parsons (1991); Moore (2011); von Wachter, Song, and Manchester (2011).

⁴Autor and Duggan (2006) discuss a number of possible explanations for the rise in DI rolls. There also exists a small body of evidence on entry responses to changes in DI benefits, wages, or local labor market conditions, including Black, Daniel, and Sanders (2002), Bratberg (1999), Campolieti (2004), Gruber (2000), and Rege, Telle, and Votruba (2009). None of these studies consider the role played by intergenerational welfare transmission.

decision in the appeal process had their case been assigned to a different judge. Our instrument picks out these complier children, whose parents are on the margin of program entry. This means we need to be cautious in extrapolating the causal effects we estimate to the population at large or to other settings. For example, the information transmitted by parents after having an appeal allowed or denied is likely to be different compared to settings where parents are on DI for other reasons (e.g., because more generous benefits induce parents to apply). Additionally, the latent demand or qualifications for DI could be higher among children whose parents are at the margin of program entry, as compared to children of inframarginal parents.

At the same time, the intergenerational link among the compliers to our instrument is relevant for policy, since reforms aimed at stemming the rise in DI will likely have the largest effect on applicants on the margin of program entry. In both Norway and the U.S., the rise in DI rolls in recent decades appears to be primarily driven by a more liberal screening of marginal applicants who are often initially denied and relatively likely to appeal (Autor and Duggan, 2006; Kostol and Mogstad, 2014). A simple simulation which makes the screening process more stringent illustrates that accounting for intergenerational effects can be important for accurate projections of post-reform participation rates and program costs. It is important to note, however, that our analysis is silent on whether the intergenerational effects we estimate are welfare improving in terms of the trade-off between costs and insurance aspects of the program.

The remainder of the paper proceeds as follows. Section 2 discusses the challenges in estimating intergenerational welfare transmission and our experimental research design. In Section 3, we describe the data, provide institutional background, and compare the DI program in Norway with that of the U.S. Section 4 presents and interprets our main findings on intergenerational welfare transmission. Section 5 explores the breadth and nature of welfare cultures in DI receipt. The final section offers some concluding remarks.

2 Identifying Intergenerational Welfare Transmission

2.1 Threats to Identification and Previous Research

In the spirit of Bertrand, Luttmer, and Mullainathan (2000), our definition of a family welfare culture is that welfare receipt in one generation causes increased participation in the next generation. This can be modeled by relating child i 's latent demand (and latent qualification) for a welfare program, P_i^{c*} , to their parent's receipt of welfare, P_i^p :

$$P_i^{c*} = \alpha^c + \beta^c P_i^p + \delta^c x_i^c + \varepsilon_i^c \quad (1)$$

where the superscripts c and p denote child and parent variables and coefficients. A child participates in the

welfare program if $P_i^{c*} > 0$. In addition to the parent's receipt of welfare, a child's participation also depends on a variety of other observable (x_i^c) and unobservable (ε_i^c) variables, such as demographic characteristics, parental characteristics, and the child's earnings capacity, health, and attitudes.

Of course, a similar equation can be written for the parent's participation decision:

$$P_i^{p*} = \alpha^p + \beta^p P_i^g + \delta^p x_i^p + \varepsilon_i^p \quad (2)$$

where the new superscript g denotes child i 's grandparent. Some of the observed x_i^p variables could also directly affect P_i^{c*} and would therefore be included in x_i^c .

A bias in the family welfare culture parameter, β^c , can arise due to unobserved factors which are correlated across generations. This becomes apparent when substituting a parent's participation resulting from equation (2) into equation (1):

$$P_i^{c*} = \alpha^c + \beta^c I(\alpha^p + \beta^p P_i^g + \delta^p x_i^p + \varepsilon_i^p > 0) + \delta^c x_i^c + \varepsilon_i^c. \quad (3)$$

where $I(\cdot)$ is the indicator function. This formulation makes clear that if $\text{corr}(\varepsilon_i^p, \varepsilon_i^c | x_i^c, x_i^p) \neq 0$, there will be a bias. For example, low earnings potential could be correlated across generations due to unobservable factors common to the parent and child, such as bad neighborhoods or low quality schools. As another example, since there is a genetic component to health, certain physical ailments could reduce work capacity within families in ways unrelated to program participation. These correlations in unobservables could incorrectly lead a researcher to believe there is a family welfare culture, when in fact the patterns are simply due to intergenerational correlations in adverse environments or poor health.

This same reasoning extends to prior generations as well. Because equation (3) is recursive, it includes a variable for the participation of a child's grandparent. If $\text{corr}(\varepsilon_i^g, \varepsilon_i^c | x_i^c, x_i^p, x_i^g) \neq 0$, this can additionally bias the family welfare culture parameter. The potential for this type of bias is suggested by studies which document multi-generational correlations in a variety of variables such as income, poverty, education, and occupation Black and Devereux (2011); Lee and Solon (2009). There is also evidence on multi-generational links in health status due to shared genes; the genetic expression of some of these conditions even skip a generation (for a review, see Bird, 2007).

Because many factors associated with welfare receipt are likely to be correlated across generations, the data demands for OLS estimation of equation (1) to yield causal evidence are high. One needs to have an exhaustive set of child and parent characteristics, as well as relevant controls for both sets of grandparents (and potentially prior generations as well). These empirical challenges have meant that existing research has largely focused on documenting the intergenerational correlations in various types of welfare use. To this

end, a number of studies have used observational data to estimate models like equation (1). For example, Bratberg, Nilsen, and Vaage (2013) provides evidence of modest, but significant, intergenerational correlation in DI receipt in Norway.⁵

While previous studies have helped researchers better describe intergenerational patterns in various types of welfare use, a causal interpretation remains elusive. As is well understood, such regressions cannot distinguish state dependence (the causal effect of welfare receipt) from that of unobserved heterogeneity (correlated unobservables across generations). There have been a few attempts to find instruments for parental welfare receipt (such as state benefit levels or local labor market conditions), include family fixed effects, or impose structural restrictions to estimate the causal intergenerational link.⁶ Pepper (2000) illustrates the difficulty in drawing credible inferences from observational data. Using a nonparametric bounds analysis, he shows that without prior information about the selection problem, the data are not informative about intergenerational welfare use. Even imposing strong assumptions or using standard instruments, he finds the bounds are wide and the point estimates are noisy and often inconsistent across specifications.

2.2 Experimental Setting and Research Design

In this subsection, we begin by reviewing key facts regarding the DI program in Norway. We then provide empirical evidence on the disability determination process, documenting in particular that the system generates random variation in DI awards. We further describe how our empirical model uses this exogenous variation to estimate the intergenerational link in DI.

The Norwegian DI program

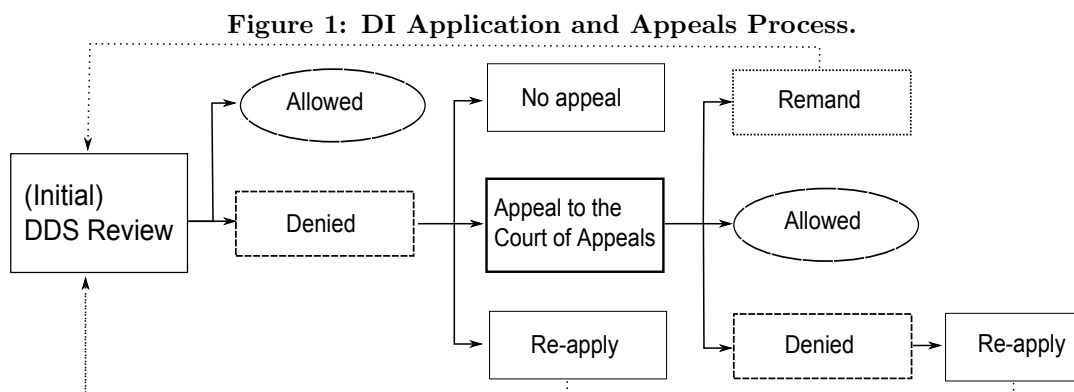
In Norway, DI benefits are designed to provide partial earnings replacements to all workers under the full retirement age who are unable to engage in substantial gainful activity because of a medically determined physical or mental impairment that has lasted for at least a year. The DI program is part of the broader Social Security System and is financed through employer- and employee-paid taxes. The level of DI benefits received is determined using a formula based on an individual's earnings history. The proportion of income that is replaced decreases as past earnings increase so that low-wage workers replace a larger fraction of their earnings than do high-wage workers. Before going on to DI, most individuals must first participate in the sick leave program which lasts at most one year and precludes full-time work. Since workers are ineligible for benefits if they can work and earn more than the substantial gainful activity threshold (about \$1,000 a

⁵For other studies of network effects in DI, see Furtado and Theodoropoulos (2012) and van Soest, Andreyeva, Kapteyn, and Smith (2011). For observational studies of other welfare programs, see e.g. Duncan, Hill, and Hoffman (1988); Solon, Corcoran, Gordon, and Laren (1988); Moffitt (1992); Antel (1992); Page and Stevens (2002) and Page (2004).

⁶See Levine and Zimmerman (1996); Gottschalk (1996); Pepper (2000); Beaulieu, Duclos, Fortin, and Rouleau (2005).

month in 2010), most workers have dropped out of the full-time labor force before applying for DI benefits.

The disability determination process involves multiple steps, as diagrammed in Figure 1. The first step is the submission of an initial application to the Social Security Administration office for the Disability Determination Stage (DDS) review. If the applicant meets the non-medical criteria (such as age and prior employment requirements), disability examiners and medical staff assess written medical evidence regarding the applicant’s ability to perform work-related activities. Examiners take into account health status, age, education, and work experience as well as the transferability of the applicant’s skills. If the disability examiner concludes that the applicant cannot be expected to engage in any substantial gainful activity, a disability award is made. Partial disability awards can also be made. Approximately 75% of claims are awarded at the DDS review. Cases that are more difficult to judge (such as mental illness and low back pain) are often denied at this step.



If the DI claim is denied at the DDS review, the individual may appeal the decision within 2 months to the Court of Appeals. About 25% of all denials are appealed. DI appeals are reviewed by Administrative Law Judges (ALJs). The ALJ must consider the application using the same criteria as the initial determination, but the applicant may present new information in writing. Judges can either allow a case, deny a case, or issue a remand (which means the case is sent back to the DDS Review stage to be re-evaluated with updated information).⁷ Approximately 15% of all claims that were appealed are allowed at the ALJ level. If the case is denied at the ALJ level, the applicant can always choose to start a new DI case by re-applying to the DDS Review stage.⁸

⁷Remands are uncommon, accounting for only 5 percent of appeal outcomes. In our baseline analysis, we code remanded cases as rejected. In a robustness check, we code remanded cases as allowed or denied based on their eventual outcome after they are reconsidered by the DDS case worker with updated information and the results are similar.

⁸Average processing time at the DDS stage is 6 months, while average processing time at the appeal stage is 4 months. Seventy-five percent of denied appellants eventually reapply, with 65 percent of these being ultimately allowed DI. If a case is denied at the ALJ level, it can also be appealed to the higher courts, but very few individuals exercise this option.

Random assignment of DI cases to judges

In Norway, the hearing of appeals is centralized in Oslo, where cases are handled for the entire country. Prior to 1998, there was only one department. Afterwards, there were four equally-sized departments; however, there is no specialization in the four departments and all judges are housed in the same building. Within each department, the assignment of a case to an Administrative Law Judge is done by the department head without knowing the content of the case, as stipulated in the rules set forth for the Administrative Law Court since its inception in 1967. The rules state that assignment should be done “by the drawing of lots.” In practice, cases are assigned on a rotating basis depending on the date they are received and the alphabetical ordering of a judge’s last name.⁹

Our setting has several attractive features: (i) the handling of cases is centralized in one location, (ii) judges do not specialize by medical condition, region of country, or other aspects of the case, (iii) the judge assesses the written evidence on the appellant’s case; there is never any personal contact between the judge and those who appeal, and (iv) an individual cannot choose an alternate judge after being assigned a judge.

A key to our design is not only that the assignment of judges is random, but also that some judges are more lenient than others. We measure judge leniency based on the average allowance rate in all other cases a judge has handled. This measure is based on all the cases a judge has ever handled, and not just those cases appearing in our estimation sample. On average, judges have handled a total of 380 cases. To construct the judge leniency measure, we calculate the leave-out mean judge allowance rate and regress this measure on fully interacted year and department dummies; this is because the randomization occurs among the pool of judges within each department. We use the residual from this regression as our judge leniency measure. This approach controls for any differences over time or across departments in the quality of applicants and the leniency of the judges.

Verifying random assignment

Table 1 empirically verifies that the hearing office complied with the random allocation procedure. This table conducts the same type of statistical test that would be done for an actual experiment to verify compliance with randomization. We find strong empirical support for the claim that the DI system in Norway randomly assigns judges to individuals who appeal their cases. The first column documents that demographic, work and health variables are highly predictive of whether an appealed case will be allowed. Column 3 examines whether our measure of judge leniency can be predicted by these same characteristics. Even though the set of characteristics are highly predictive of case outcomes, they are not statistically related to the leniency of

⁹We verified these rules with the current Head of the Administrative Law Court, Knut Brofoss. The rules are explained in “Veileder for Saksbehandlingen i Trygderetten” (Guidelines for Processing Cases in the Court of Appeals).

the judge assigned to a case: none of the 19 variables are statistically significant at the 5% significance level and the variables are not jointly significant either.¹⁰ In fact, the point estimates are close to zero, and taken together, the variables explain only 0.35 percent of the variation in our measure of judge leniency. Note in particular the insignificance of the disorder variables. This is consistent with the lack of specialization by type of disability in Norway, something which is not true in many other countries.

Table 1: Testing for Random Assignment of Cases to Judges.

	Dependent Variable			
	Case Allowed		Judge Leniency	
	coeff.	s.e.	coeff.	s.e.
Age	0.00539***	(0.00088)	0.00036*	(0.00020)
Female	0.01088	(0.00966)	0.00022	(0.00189)
Married	0.00419	(0.00760)	0.0013	(0.00191)
Foreign born	-0.02713***	(0.01140)	0.00094	(0.00246)
Less than high school	-0.01670***	(0.00704)	-0.00027	(0.00175)
High school degree	0.01317*	(0.00700)	0.00041	(0.00143)
Some college	0.02282	(0.01613)	-0.00073	(0.00337)
College graduate	-0.10339***	(0.01991)	0.00389	(0.00949)
One child	-0.0052	(0.00878)	-0.00097	(0.00200)
Two children	-0.01593	(0.01322)	0.00103	(0.00164)
Three or more children	-0.03559***	(0.01461)	0.00319	(0.00214)
Average indexed earnings	0.00000***	(0.00000)	0.00000	(0.00000)
Experience	0.00520***	(0.00086)	0.00001	(0.00022)
Mental disorders	0.03572***	(0.01054)	0.00005	(0.00384)
Musculoskeletal disorders	0.00263	(0.00861)	0.0018	(0.00256)
Circulatory system	0.01271	(0.02981)	-0.00219	(0.00427)
Respiratory system	0.01453	(0.02338)	0.00634	(0.00423)
Nervous system	0.06380**	(0.03162)	0.00422	(0.00434)
Endocrine diseases	0.00614	(0.02578)	-0.00088	(0.00466)
F-statistic for joint significance	9.25		.77	
[p-value]	[.001]		[.730]	
N	14,722		14,722	
R-squared	.0155		.0035	

***p<.01, **p<.05, *p<.10. Standard errors (in parentheses) are clustered at the judge level.

Notes: Baseline estimation sample, consisting of parents who appeal an initially denied DI claim during the period 1989-2005 (see Section 3 for further details). There are 79 different judges. The judge leniency variable is constructed by calculating the leave-out mean judge allowance rate for all cases a judge has handled (not just those in the baseline estimation sample), regressing this measure on fully interacted year and department dummies, and using the residual from this regression as the variable. Columns 1 and 3 display OLS estimates from separate regressions of whether a case is allowed or judge leniency, respectively, on appellant characteristics. F-statistics are obtained from OLS estimation on the combined set of applicant characteristics. All regressions include fully interacted year and department dummies. Characteristics of appellants are measured prior to the appeal. Number of children is the number under age 18, average indexed earnings is mean earnings for the last ten years prior to appeal and experience is number of years with positive earnings over this ten year period.

A natural question is why some judges are more lenient than others. While we do not have detailed

¹⁰The coefficient on age, while close to zero, is statistically significant at the 10% level. Given the number of covariates we consider, this is not surprising, since the probability of observing one p-value at this level by chance alone is large.

characteristics of the judges, we do know the number of cases they have handled. Whereas experienced judges appear to be slightly less lenient, experience accounts for only a small fraction of the total variation in allowance rates across judges (see Appendix Figure A.1). Other unobserved factors must be driving the underlying variation. It is important to recognize that as long as judges are randomly assigned, it does not matter why some judges are more lenient than others.

Instrument and empirical model

We use variation in DI allowance generated from the random assignment of appeal judges as an instrument to estimate the intergenerational link in DI. We estimate judge leniency by taking the average allowance rate in all other cases a judge has ever handled (not just the cases in our estimation sample), adjusted for year and department effects, as we did for Table 1.¹¹ As we document below, some judges are systematically more lenient than others, which gives exogenous variation in the probability a parent is allowed DI in the appeals process.

Our baseline empirical model can be described by the following two-equation system:

$$A_i = \alpha + \gamma z_i + \theta x_i + v_i \tag{4}$$

$$P_i = \mu + \beta A_i + \lambda x_i + u_i \tag{5}$$

where z_i denotes a judge’s leniency, A_i is an indicator for whether the parent is allowed DI in the appeal process, P_i is an indicator variable for whether the child subsequently participates in DI, and x_i is a vector of control variables. We perform 2SLS with equation (4) as the first stage and equation (5) as the second stage, with the goal of consistently estimating the parameter β . We think of this parameter as a measure of family welfare culture, giving the effect of a parent being allowed DI because of a lenient judge on their adult child’s DI participation. Our estimate captures any effect which operates through whether the parent is allowed DI in the appeal process, including participation in DI, subsequent reapplications to the DI program or any other causal change in parental behavior. We can also estimate the reduced form effect by directly regressing P_i on z_i and x_i .

At the outset, it is important to be precise about the causal effect being estimated. Influential work by Imbens and Angrist (1994) has clarified the interpretation of 2SLS estimates as local average treatment effects (LATE) when β is a random coefficient. Applied to our setting, this means the welfare culture parameter pertains to children whose parents could have received a different allowance decision in the appeal process had their case been assigned to a different judge. As discussed in greater detail later, this suggests

¹¹Although the instrument is pre-estimated, there is no need to adjust the standard errors of the IV estimates; such adjustments are necessary with generated regressors but not with generated instruments.

due caution in extrapolating the causal effects we estimate to the population at large or to other settings. For example, the information transmitted by parents after having an appeal allowed or denied is likely to be different compared to settings where parents are on DI for other reasons (e.g., because more generous benefits induce parents to apply). Additionally, the latent demand or qualifications for DI could be higher among children whose parents are at the margin of program entry, as compared to children of inframarginal parents.

3 Data and Background

3.1 Data and Sample Restrictions

Our analysis employs several data sources that we can link through unique identifiers for each individual. Information on DI benefits comes from social security registers that contain complete records for all individuals who entered the DI program during the period 1967-2010. The data set includes information on the individual's work history and medical diagnosis, the month when DI was awarded (or denied), and the level of DI benefits received. We link this information with administrative data from the hearing office on all appeals from 1989 to 2011. The data set contains information on dates of appeal and decision, the outcome of the appeal, and unique identifiers for both judges and applicants. We merge these data sets with administrative registers provided by Statistics Norway, using a rich longitudinal database that covers every resident from 1967 to 2010. For each year, it contains individual demographic information (including sex, age, and number of children), socio-economic data (such as years of education and earnings), and geographical identifiers. The data contains unique identifiers that allow us to match parents to their children, as well as spouses and siblings to each other. We can further match neighbors to each other using street addresses. The coverage and reliability of Norwegian registry data are rated as exceptional in international quality assessments (see Atkinson, Rainwater, and Smeeding 1995).

Our empirical analysis considers children of parents who appeal an initially denied DI claim.¹² Following Maestas, Mullen, and Strand (2013) and French and Song (2013), our baseline estimation excludes observations for which the assigned appeal judge has handled few cases (less than ten during the period 1989 to 2011). The reason for this sample restriction is to reduce the noise in our instrument. We further refine the sample to be appropriate for studying intergenerational transmission of DI receipt. We begin by restricting the sample to children whose parent's appeal decision was made during the period 1989 to 2005. This sample restriction allows us to observe the behavior of children for at least five years after the appeal

¹²Some parents have several denied DI claims over the period we consider. In such cases, we restrict our sample to the parent's first denied DI claim.

decision of the parent. We further exclude children whose parent were older than 55 years at the time he or she appealed. The reason for this age restriction is to avoid program substitution between DI and early retirement schemes.

In our main analysis, we restrict the sample to children who are age-eligible for DI (at least 18 years old) at the time of the parent’s appeal decision. This age restriction allows us to observe participation behavior over time for a sizeable sample of children. The baseline sample consists of 14,722 parent-child observations and 79 different judges; our sample includes roughly two children over the age of 18 per parent. One implication of the age restriction is that the baseline sample will be comprised of older children as compared to the unrestricted sample of appellants. Appendix Figure A.2 displays the age distribution of parents who appeal and the age distribution of their children. Because few parents with young children apply for DI, the baseline sample includes the typical parent-child links. In Section 5, we will nevertheless explore the impact of parental DI participation on an alternative, smaller sample of children who are under 18 at the time of the parent’s appeal decision.

In Table 2, we document the key characteristics of the sample of parents who apply for DI and our baseline sample of parents who appeal an initially denied DI claim. The parents who appeal are on average more likely to be female, less educated and foreign born, and have lower prior earnings and less work experience compared to the group of initial applicants. Sixty-five percent of applicants claim mental or musculoskeletal disorders, a percentage that rises to 73 percent for appellants. The children of parents who appeal tend to be less educated, but actually have slightly higher prior earnings compared to children of parents who initially apply for DI. In the time span we observe, the children of parents who appeal are slightly more likely to be DI recipients compared to children of parents who initially apply for DI (8 percent versus 7 percent).¹³ While every child is observed for at least five years, some children will be observed for up to 21 years; on average, a child is observed for 11 years.

3.2 Institutional Background

There are a number of similarities and a few key differences between the DI systems in the U.S. and in Norway (see Autor and Duggan, 2006; Kostol and Mogstad, 2014). In both countries, DI is one of the largest transfer programs. However, the incidence of receipt of DI benefits is lower in the U.S. than in Norway. Figure 2 shows this distinction by displaying the evolution of DI in the two countries. Whereas the rate of DI receipt in a given year is consistently higher in Norway than in the U.S., the time trends are quite

¹³By way of comparison, the rate of DI receipt is equal to 3 percent for a comparable set of children whose parents never applied for DI. To create a comparable set, we matched on the covariates appearing in Panel B of Table 2 (except for type of disability, which is not available in both datasets).

similar.¹⁴ From 1961 to 2012, the rate of receipt increased from 2.2 to 9.7 percent in Norway and from 0.8

Table 2: Descriptive Statistics.

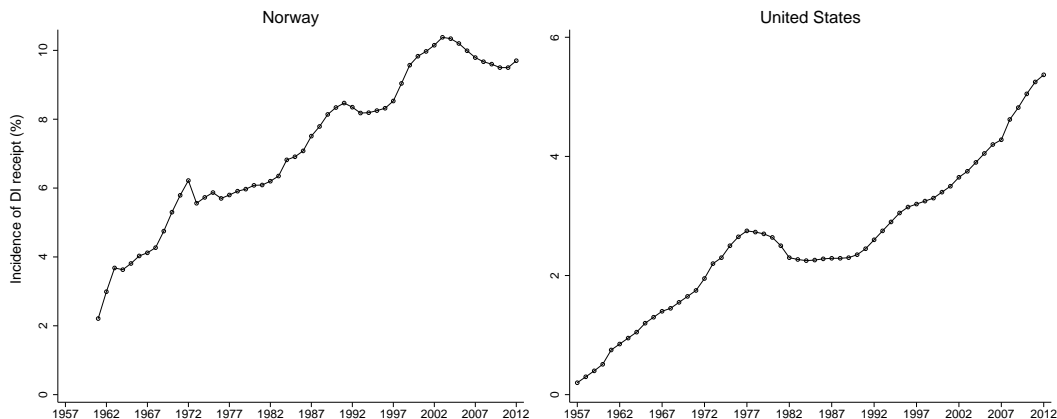
Characteristic	DI applicants		DI appellants	
	Mean	Std. Dev.	Mean	Std. Dev.
A. Parents				
Age (time of decision)	47.28	[7.09]	49.19	[4.36]
Female	0.65	[0.48]	0.74	[0.44]
Married	0.62	[0.48]	0.68	[0.46]
Foreign born	0.09	[0.28]	0.18	[0.38]
Less than high school	0.44	[0.50]	0.55	[0.50]
High school degree	0.44	[0.50]	0.38	[0.48]
Any college	0.12	[0.33]	0.07	[0.26]
Children below age 18 living at home	0.41	[0.49]	0.43	[0.49]
Previous earnings (\$), 1-10 years prior to decision	29,721	[22,052]	20,681	[19,037]
Years of work, 1-10 years prior to decision	7.94	[3.06]	6.78	[3.61]
Mental disorders	0.24	[0.43]	0.21	[0.41]
Musculoskeletal disorders	0.41	[0.49]	0.52	[0.50]
Circulatory system	0.06	[0.25]	0.04	[0.19]
Respiratory system	0.03	[0.16]	0.03	[0.17]
Nervous system	0.06	[0.23]	0.03	[0.17]
Endocrine diseases	0.02	[0.13]	0.04	[0.19]
DI allowed	0.75	[0.43]	0.12	[0.32]
Number of parents	98,206		7,331	
B. Children				
Age (time of decision)	25.33	[4.63]	24.98	[4.63]
Female	0.43	[0.49]	0.49	[0.50]
Married	0.14	[0.35]	0.16	[0.37]
Foreign born	0.17	[0.38]	0.13	[0.33]
Less than high school	0.49	[0.50]	0.52	[0.50]
High school degree	0.37	[0.48]	0.37	[0.48]
Any college	0.14	[0.35]	0.12	[0.32]
Children below age 18 living at home	0.4	[0.49]	0.31	[0.46]
Previous earnings (\$), 1-5 years prior to decision	19,326	[20,776]	20,682	[20,680]
Years of work, 1-5 years prior to decision	3.45	[1.94]	3.73	[1.68]
DI recipient 5 years after decision	0.03	[0.17]	0.03	[0.16]
DI recipient any time after decision	0.07	[0.25]	0.08	[0.27]
Number of children	195,223		14,722	

Notes: Sample of parents and children for applicants during the period 1992-2005 and appellants during the period 1989-2005. In both samples parents are restricted to be at most age 55 and their children to be aged 18 and above at the time of decision (at the application step or the appeal step). Previous earnings and years of work are measured the year before appeal in the DI appellant sample and the year before decision in the DI applicant sample. Nominal values are deflated to 2005 and represented in US dollars using the average exchange rate NOK/\$ = 6. Unless otherwise stated, all parent and child characteristics are measured the year before parental application/appeal.

¹⁴The cross-country difference in DI coverage is unlikely to explain the entire discrepancy in the incidence of DI: although virtually all non-elderly adults are covered in Norway, more than 80 percent of all non-elderly adults are covered in the U.S. The remaining difference could be a function of underlying differences in screening stringency, the generosity of the programs or the frequency with which people apply for disability benefits. Milligan and Wise (2011) argue that differences in health are unlikely to explain much of the observed differences in DI rates across developed countries.

to 5.4 percent in the U.S. While Norway’s rate has leveled off at about 10 percent in recent years, the U.S. DI rate continues to rise and is projected to exceed 7 percent by 2018 (Burkhauser and Daly, 2012).

Figure 2: Trends in DI Receipt in Norway and the U.S.



Notes: U.S. trends based on Autor and Duggan (2006) for 1957-2005 and SSA Office of the Chief Actuary for 2006-2012. Norwegian trends based on SSA Statistical Supplements. Incidence of DI receipt is defined as the percent of the relevant adult population receiving DI benefits (age 18-67 in Norway; age 25-64 in the US).

In both countries, the expansion of the DI rolls in recent decades appears to be driven by the liberalization of the screening process, which led to a rapid increase in the share of DI recipients suffering from difficult-to-verify disorders such as mental illness and musculoskeletal disease.¹⁵ Because these are early-onset disorders with low mortality at young ages, DI recipients with such diagnoses tend to participate in the program for relatively long periods. As a result, the DI exit rates in both countries have decreased in the last few decades, with progressively fewer DI recipients reaching retirement age or dying in a given year (see Appendix Figures A.3 and A.4).

There are a few noticeable differences between the two countries. DI recipients in Norway tend to be older and have slightly higher earnings prior to a disability award. One possible explanation for this is that the U.S. SSDI program is less generous.¹⁶ The differences in characteristics are, however, less pronounced than one might expect. For instance, almost 60 percent of DI recipients suffer from difficult-to-verify disorders (mental illness and musculoskeletal disorders) in both the U.S. and Norway (see Appendix Table A.1).

Another difference is that the appeal process plays a more important role in the U.S. than in Norway. While 48 percent of the initially rejected applicants appeal in the U.S. (French and Song (2013)), only

¹⁵See Autor and Duggan (2006) for a discussion of this phenomenon. In the U.S., the 1984 congressional reforms shifted the focus of screening from medical to functional criteria. In Norway, the medical eligibility criteria were relaxed earlier and more gradually.

¹⁶For a typical DI recipient in Norway, Kostol and Mogstad (2014) calculate the replacement rate would be 31 percent according to U.S. program rules and 58 percent according to Norwegian program rules. Factoring in health insurance coverage increases the effective replacement rate to over 50 percent in the U.S. In Norway, all citizens are eligible for health insurance through the Social Insurance system.

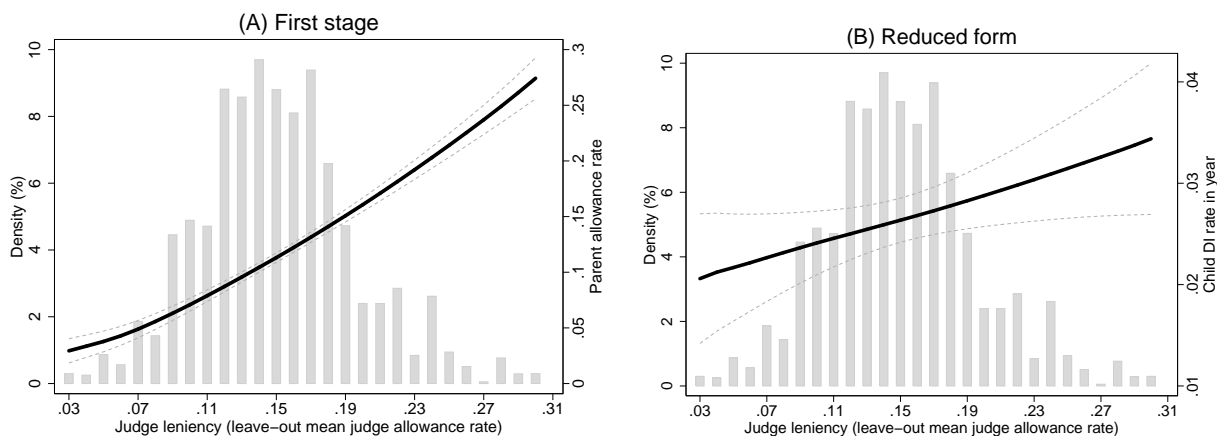
25 percent of the initially rejected appeal in Norway. Appendix Table A.1 compares the characteristics of individuals who apply for DI and those who appeal an initially denied DI claim in the two countries. In both the U.S. and Norway, appellants are more likely to be younger, less connected to the labor market, and more likely to suffer from difficult-to-verify disorders, as compared to the the initial group of applicants. This suggests that in both countries the marginal applicants are often initially denied, and they are relatively likely to appeal.

4 Evidence on Intergenerational Welfare Transmission

4.1 Graphical Evidence

We begin our presentation of results by providing a graphical representation of the IV approach in Figure 3. In the background of each graph is a histogram for the density of judge leniency, which captures the average judge allowance rate in the other cases a judge has handled. We note the judge leniency measure is calculated from all cases the judge has ever handled, not just the cases in our estimation sample. On average, each judge has handled a total of 380 cases. The mean of the leniency variable is .15 with a standard deviation of .06. The histogram reveals a wide spread in judge leniency, with approximately 22% of cases allowed by a judge at the 90th percentile compared to approximately 9% at the 10th percentile.

Figure 3: Effect of Judge Leniency on Parents (First Stage) and Children (Reduced Form).



Notes: Baseline sample, consisting of parents who appeal an initially denied DI claim during the period 1989-2005 (see Section 3 for further details). There are 14,722 individual observations and 79 different judges. Panel (A): Solid line is a local linear regression of parental DI allowance on judge leniency. Panel (B): Solid line is a local linear regression of child DI receipt on their parent’s judge leniency measure. All regressions include fully interacted year and department dummies. The histogram of judge leniency is shown in the background of both figures (top and bottom 1% excluded from the graph). Dashed lines represent 90 percent confidence intervals.

Panel A shows the effect of judge leniency on a parent’s allowance rate. The graph is a flexible analog to the first stage equation (4), where we plot a local linear regression of actual parental allowance against

judge leniency. The parental allowance rate is monotonically increasing in our leniency measure, and is close to linear. A 10 percentage point increase in the judge's allowance rate in other cases is associated with an approximately 9 percentage point increase in the probability the parent's case is allowed. Panel B plots the reduced form effect of a parent's judge leniency measure against their child's DI participation, again using a local linear regression. The child's DI rate is monotonically increasing in the leniency measure as well. Approximately two and a half percent of children whose parents had a relatively strict judge (leniency measure = .09, the 10th percentile) are predicted to participate in DI five years later. This can be contrasted with roughly three percent of children whose parents had a relatively lenient judge (leniency measure = .22, the 90th percentile).

4.2 Regression Estimates

We now turn to a regression based analysis. Column 1 in Table 3 reports first stage estimates which regress a dummy variable for whether a parent is allowed DI at the appeal stage on our judge leniency measure. We include fully interacted year and department dummies in the first column, but otherwise include no other controls. The coefficient implies that when a judge's allowance rate in the other cases he has handled goes up by 1 percentage point, the probability a parent will be allowed DI by that judge increases by 0.91 percentage points. This effect is not statistically different from one.

Panel A reports results for whether the child participates in DI within 5 years after the parent's appeal decision. Column 2 reports the reduced form estimate of a parent's judge leniency measure for this child outcome. The estimate of .055 implies that when judge leniency for a parent rises by 10 percentage points, a child's DI participation will rise by roughly one-half of a percentage point. This is a sizeable effect compared to the 3 percent average DI participation rate within five years for this sample. Column 5 takes the reduced form estimate of column 2 and divides it by the first stage estimate in column 1. Since the first stage is close to one, the reduced form and the IV estimates are very similar.

Panel B performs a similar exercise, but now looks at whether the child has ever been on DI after the parent's appeal decision. While every child is observed for at least five years after their parent's appeal decision, in this second panel some children will be observed for up to 21 years and, on average, the children are observed for 11 years. The unbalanced nature of this second panel affects the interpretation of the estimates, but it should not affect their validity given the nature of our instrument. Figure 4 complements Table 3 by showing IV estimates for the intergenerational transmission of DI receipt over time for a balanced panel. The estimates correspond to those in Table 3, except the graph restricts the sample to children observed for at least 10 years after their parent's appeal decision.¹⁷ Both Table 3 and Figure 4 suggest the

¹⁷The first stage estimate for this sample is 1.006 with a standard error of 0.146.

long-run effects of a parent getting on to DI are roughly twice as large as the short-run effects. For example, ten years after the court decision, Figure 4 reveals that the causal effect of a parent being allowed DI is a 12 percentage point increase in a child’s DI take up. These findings suggest that a parent’s experience with the DI system is not merely changing the timing of when their children participate in DI, at least over the time period covered by our data.

Table 3: Estimates of Intergenerational Welfare Transmission.

	<i>Child on DI 5 years after parent’s appeal decision</i>		<i>Child ever on DI after parent’s appeal decision</i>		
	First stage	Reduced form	IV	Reduced form	IV
A. No additional controls					
Parent’s judge leniency	0.909*** (0.112)	0.055*** (0.020)		0.107*** (0.030)	
Parent allowed DI			0.061*** (0.022)		0.118*** (0.033)
B. With additional controls					
Parent’s judge leniency	0.869*** (0.108)	0.052** (0.020)		0.101*** (0.027)	
Parent allowed DI			0.060*** (0.023)		0.116*** (0.032)
Dependent mean	0.12	0.03		0.08	

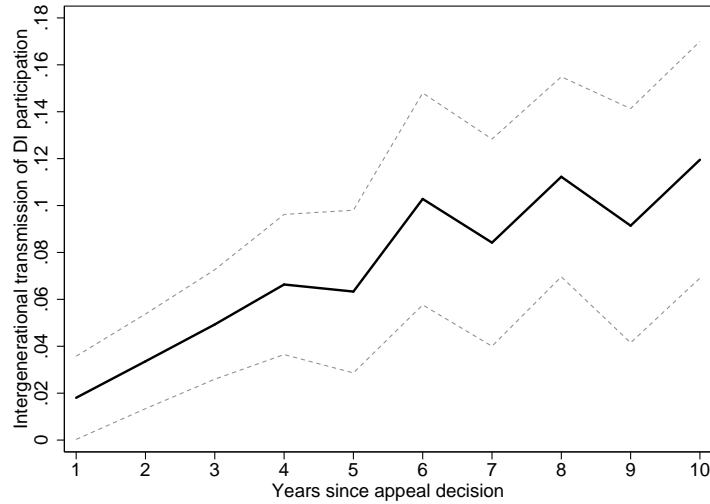
***p<.01, **p<.05, *p<.10. Standard errors (in parentheses) are clustered at the judge level.

Notes: Baseline sample of 14,722 child-parent observations, restricted to parents who appeal an initially denied DI claim during the period 1989-2005 (see Section 3 for further details). There are 79 different judges. All regressions include fully interacted year and department dummies. Specifications with additional controls include a linear term for average indexed earnings and dummy variables for month of appeal, county of residence, age of parent and child, gender of parent and child, foreign born, marital status, number of children, education, labor market experience, and a number of medical diagnoses. The control variables are measured prior to the appeal. Number of children is the number under age 18, average indexed earnings is mean earnings for the last ten years prior to appeal and experience is number of years with positive earnings over this ten year period.

This rising trend in the estimates captures both the effect of elapsed time since parental allowance as well as the effect of children getting older. If the causal effect is larger for older children, then the aging of children over time could be the underlying reason for the trend. To explore this possibility, in Appendix Table A.5 we reweight individual observations so that the distribution of ages is the same in each year and centered around a mean age of 30. Holding the age distribution constant in this way, we then re-estimate the effects over time. The estimates are remarkably similar, with the reweighted estimates growing substantially over time as before, indicating that elapsed time since parental allowance is a key reason for how the intergenerational transmission evolves.

Lastly, we shift attention to how a parent’s DI receipt affects the probability that their children subsequently apply for DI. Appendix Figure A.5 shows IV estimates for child DI application over time based on

Figure 4: Estimates of Intergenerational Transmission over Time.



Notes: Baseline sample restricted to parents who appeal an initially denied DI claim during the period 1989-2000, so as to have a balanced 10 year sample. There are 9,062 individual observations and 50 different judges. The figure displays separate IV estimates of intergenerational transmission 1 to 10 years after the parent’s appeal decision. The specifications mirror column 3 of Panel A in Table 3. Dashed lines represent 90 percent confidence intervals (clustered at the judge level).

the ten-year balanced panel used in 4. These results mirror closely the estimates for DI participation. The effect on DI application grows substantially over time. Ten years after the court decision, the causal effect of a parent being allowed DI at the appeal stage is a 14 percentage point increase in a child’s DI application rate. Given the qualitative similarity in the estimates when using child application versus child participation as the left hand side variable, we focus on children’s DI participation in the remainder of the paper.

4.3 Internal Validity

In order for judge leniency to be a valid instrument, appellants’ assignment to judges must be uncorrelated with case characteristics. Table 1 provided strong empirical support for the claim that the DI system in Norway randomly assigns appeal judges within each department and year. As a second test, Panel B of Table 3 explore what happens if a large set of control variables are added to the baseline regressions. If judges are randomly assigned, the addition of these control variables should not significantly change the estimates, as both parental and child characteristics should be uncorrelated with judge leniency. As expected, the coefficients do not change appreciably. As a final test of randomization, we examine whether the likelihood of children receiving sickness pay prior to the parents’ appeal is correlated with judge leniency. Before going onto DI, individuals usually participate in the sickness program; correlation between our instrument and children’s pre-determined participation rate in this program would therefore raise concerns about compliance with the random allocation procedure. It is reassuring to find that child participation in the sickness program

is not statistically related to the leniency of the judge assigned to their parent’s case.¹⁸

While random assignment of cases to judges is sufficient for a causal interpretation of the reduced form estimates, the IV estimates require two additional assumptions. The first is that the leniency of the parent’s judge affects the child’s DI participation only through the parent’s allowance decision, and not directly in any other way. One attractive feature of the process in Norway makes this exclusion restriction likely to hold: the appeal is presented in writing, so there is never any personal contact between the judge and those who appeal. What parents and children observe is the allowance or denial decision of the judge.

A possible caveat is that appeal processing time could differ systematically by the leniency of the judge (see e.g. Autor, Maestas, Mullen, and Strand (2011)) and that this could directly affect a child’s decision to apply for DI. To examine this, we calculated a judge’s average processing time based on the residual average processing time in the other cases a judge has handled after controlling for a fully interacted set of time and department dummies in a regression. It is reassuring to find that our instrument, judge leniency, and judge processing time are virtually uncorrelated. Moreover, the second row of Table 4 shows that the IV estimates do not change appreciably if we control for a judge’s average processing time (an exogenous variable since judges are randomly assigned) in the first and second stages.

The final assumption needed for a causal interpretation of the IV estimates is monotonicity of judges’ appeal decisions. In our setting, the monotonicity assumption is that cases allowed by a strict judge would also have been allowed by a more lenient judge, and similarly that cases denied by a lenient judge would also have been denied by a stricter judge. One testable implication of the monotonicity assumption is that the first stage estimates should be non-negative for all subsamples. Appendix Table A.4 provides separate first stage estimates based on characteristics of the parent and the child. These estimates are consistently positive and sizeable, in line with the monotonicity assumption.

Lastly, Table 4 reports the results from several specification checks, all of which support our main findings. In specification C, we limit the sample to the period when there was just one department, rather than four departments handling appeals. While the standard errors go up somewhat, the results are similar. Specifications D and E show the results are robust to adding in fully interacted year, month and department dummies or excluding parents who die. In our baseline analysis, we excluded judges who handle less than 10 cases. Specifications F and G demonstrate that including these judges does not change the estimates appreciably, and neither does excluding judges who handle less than 50 cases. Specification H considers an alternative handling of remanded cases. In our baseline analysis, we code a remanded case as rejected (see footnote 7). If we instead code remanded cases as allowed or denied based on its eventual outcome after it

¹⁸The regression coefficient of parental judge leniency on a child’s participation in the sickness program is 0.004 (s.e. = 0.05). This point estimate is small compared to the sample mean: 24 percent of children had received sickness pay at some point prior to their parent’s appeal.

is reconsidered by the DDS case worker with updated information, the results are quite similar. The final specification drops appeals where the claim was made after January 1, 2004 since the DI system was reformed starting that year. The estimates do not change appreciably.

Table 4: Specification Checks for Intergenerational Welfare Transmission Estimates.

Specification	<i>Child on DI 5 years after parent's appeal decision</i>			
	First stage	Reduced form	IV	N
A. Baseline specification	0.869*** (0.108)	0.052** (0.020)	0.060*** (0.023)	14,722
B. With judge ave. processing time	0.851*** (0.103)	0.050** (0.021)	0.059** (0.023)	14,722
C. One Department (pre-1998)	1.000*** (0.143)	0.061** (0.025)	0.061** (0.028)	5,567
D. Month-department controls	0.777*** (0.122)	0.049** (0.022)	0.063** (0.029)	14,722
E. Exclude parents who die	0.867*** (0.108)	0.060*** (0.021)	0.070*** (0.024)	14,314
F. Include judges < 10 cases	0.858*** (0.109)	0.051** (0.020)	0.060*** (0.023)	14,726
G. Exclude judges < 50 cases	0.950*** (0.102)	0.054** (0.022)	0.056** (0.022)	14,587
H. Alternative coding of remand	0.808*** (0.100)	0.053*** (0.018)	0.066*** (0.022)	14,722
I. Drop cases after 1/1/2004 reform	0.891*** (0.107)	0.052** (0.021)	0.058** (0.023)	14,474

***p<.01, **p<.05, *p<.10. Standard errors (in parentheses) are clustered at the judge level.

Notes: Specifications mirror the baseline results with additional controls reported in panel B of Table 3.

4.4 Interpreting the IV Estimates

It is important to emphasize the local nature of our results. Our IV estimates represent a local average treatment effect (LATE) for children whose parents could have received a different allowance decision had their case been assigned to a different judge. To better understand this LATE, we take several steps.

Number of complier children and their characteristics

We begin by calculating the number of children whose parents are always takers, never takers and compliers in our sample. Appendix B provides details for these calculations. Compliance types are usually defined in the context of binary instruments. However, the approach of Imbens and Rubin (1997) and Abadie (2003) extends naturally to our setting with a continuous instrument, by looking at the allowance rates for parents who are assigned to the “most lenient” and the “strictest” judges. Parental compliers are appellants who would have received a different allowance decision had their case been assigned to the most lenient judge

instead of the strictest judge. We estimate that children of compliers make up approximately 25 percent of our sample. Because of monotonicity, the share of parents that would be allowed DI regardless of the judge assigned to their case is given by the probability of allowance for the strictest judge. The children of these always takers make up only a few percent of the sample. By comparison, more than 70 percent of our sample are children of never takers who would not be allowed DI no matter which judge was assigned to their case.

We characterize compliers by observable characteristics in Appendix Table A.4. As explained in Abadie (2003), these characteristics can be recovered by calculating the fraction of compliers in different subsamples. The most distinctive feature of the compliers is their family background: 65 percent of complier children have parents with low education, while their fraction in the entire sample is only 56 percent. Parental compliers are also more likely to have difficult-to-verify disorders as compared to other types of parents who appeal.

Potential participation rates of complier children

The IV estimates reveal the probability a complier child has ever been on DI after their parent's appeal increases by 12 percentage points if the parent is allowed DI. A natural question is: how many complier children would have been on DI if their parents had been denied DI? As shown in Appendix B, we can recover this potential outcome by combining (i) our estimates of the shares of never takers and compliers with (ii) estimates of the mean child participation rates of children whose parents were not allowed with the most lenient or strictest judges. We find that roughly 3 percent of the complier children could have ever been on DI after their parent's appeal if their parent had been denied DI, a fraction which is lower than the 8 percent observed for all children of appellants in our sample.

Labor and educational outcomes of children

In Table 5, we explore the labor and educational outcomes for complier children. Consistent with the intergenerational impact on children's use of DI, we find that parental DI allowance decreases the probability that a child will be employed or pursue higher education. Examining child outcomes five years after their parent's appeal, Table 5 shows that a parent's DI receipt causes employment to drop by 14 percentage points. While we do not estimate the drop in full-time work or college completion with the same precision, both estimates suggest a sizeable drop in these child outcomes as well. Taken together, the estimated effects indicate that parental DI allowance induces welfare participation among children who otherwise would have been inclined to work or invest in education. This finding is important for accurate projections of the overall economic consequences of tightening the screening process.

Table 5: Effect of Parent’s DI Allowance on Child Labor and Educational Outcomes.

5 years after parent’s appeal decision

Dependent variable	Reduced form	IV	Dep. mean
A. DI	0.052** (0.020)	0.060*** (0.023)	0.03
B. Any employment	-0.119** (0.055)	-0.137** (0.065)	0.58
C. Full-time work	-0.065 (0.079)	-0.075 (0.090)	0.42
D. College degree	-0.079 (0.060)	-0.091 (0.069)	0.25

***p<.01, **p<.05, *p<.10. Standard errors (in parentheses) are clustered at the judge level.

Notes: Baseline sample of 14,722 observations (see Table 3). There are 79 different judges. Specifications mirror the baseline results with additional controls reported in panel B of Table 3. Any employment is defined as working more than 4 hours a week, full-time work as more than 30 hours a week, and college degree as having completed college by 2010. Labor outcomes are measured five years after parent’s appeal decision.

Extrapolation of LATE and comparison with OLS

The welfare culture parameter we estimate is specific to children whose parents would have received a different allowance decision in the appeal process had their case been assigned to a different judge. Our instrument picks out these complier children, whose parents are on the margin of program entry. This suggests due caution in extrapolating the causal effects we estimate to the population at large or to other settings. Additionally, we need to be cautious in comparing the local effects for children of complier parents to OLS estimates of equation (5). This point has been emphasized in previous work that use a similar identification approach based on quasi-random assignment of judges (or examiners) in other contexts.¹⁹

When estimating equation (5) using OLS, we find a very weak association between child DI participation and a parent’s DI allowance. Using the same sample as in Table 3, the OLS estimate of parental allowance on whether a child is ever on DI is .01 (s.e.=.01), a number which is close to zero and considerably smaller than our IV estimate. The OLS estimate can differ from the IV estimate for at least two reasons. The first is heterogeneity in effects and the second is selection bias due to correlated unobservables. In our setting where all parents have chosen to apply for DI and appeal the initial rejection, it is difficult to predict the nature of the heterogeneity or sign the direction of the bias. For example, if genetic components to health are important in judges’ decision to allow DI at the appeal stage, then we would expect upward biased OLS estimates; on the other hand, OLS estimates would be biased downwards if judges are inclined to award DI in cases where “random” diseases or accidents happen to healthy people.

To better understand the relatively low value of the OLS estimate, we take several steps. We begin by

¹⁹For example, see the discussions in French and Song (2013) and Maestas, Mullen, and Strand (2013) about why estimates of the effect of DI allowance on labor supply from IV might be equal to or exceed those from OLS.

testing for heterogeneity in welfare transmission, exploiting that constant-effects models with a multivalued instrument and a binary endogenous regressor are over-identified. Specifically, we construct a set of dummy instruments for ten equally spaced intervals of the support of the underlying multivalued instrument. The 2SLS estimator using this set of dummy instruments is the efficient linear combination of all the just-identified IV estimators generated by these instruments one at a time. Under the null hypothesis of constant effects, this 2SLS estimator should not be significantly different from our baseline IV estimator which uses the linear instrument. When performing this test, we can reject the null hypothesis of homogenous effects of parental DI allowance at conventional levels of significance. As a result, the difference between the OLS and IV estimates cannot be attributed to selection bias only.

Next, we decompose the OLS estimate (see Appendix B for details). The OLS estimand can be expressed as the difference in the observed participation rate of children whose parents are allowed DI (treated) and those whose parents are denied DI (untreated). As usual, the treated consists of always takers and treated compliers, whereas the untreated consists of never takers and untreated compliers. The OLS estimand can therefore be written as the weighted average of childrens' potential participation rates if their parents were allowed DI across the always takers and treated compliers, minus the weighted average of childrens' potential participation rates if their parents were denied DI across the never takers and untreated compliers. Both the potential outcomes and the weights are possible to compute directly from data; by calculating these quantities, we can therefore learn more about why the IV estimate differs from the OLS estimate in our setting with heterogenous effects.

Our calculations reveal that a treated child's potential participation rate if their parent is allowed DI is similar in magnitude to an untreated child's potential participation rate if their parent is denied DI (they are both close to 8%). As a result, the implied OLS estimate is close to zero, mirroring closely the OLS estimate we obtain when we regress child DI participation on parental DI allowance directly. Our calculations also reveal that a child's potential participation if their parent is allowed DI is smaller for always takers than for compliers, whereas a child's potential participation if the parent is denied DI is larger for never takers than for compliers. Taken together, this pattern of heterogeneity generates an OLS estimate that is close to zero, especially since children of never takers make up most of the untreated group (about 81 percent).

These findings are consistent with the notion that the beliefs or attitudes of complier children are more sensitive to the outcome of the parent's appeal compared to always takers and never takers. Consider first the comparison between never takers and untreated compliers. A key difference between these two groups is that compliers are on the margin of being allowed DI, while never takers have relatively clear-cut cases (they will not be allowed DI no matter which judge they are assigned to). It is therefore plausible that DI receipt by a parent is more informative and salient for a complier child's beliefs. For example, a complier

child whose parent is denied DI because of assignment to a strict judge may infer that applying for DI is not worthwhile for marginal cases. By comparison, seeing that clear-cut cases are not allowed at the appeal stage might lead to only small changes in a child's beliefs, attitudes and application behavior. A similar logic applies to the comparison between always takers and treated compliers. Always takers will be allowed DI even if they get the strictest judge, so observing these clear-cut cases being allowed at the appeal stage might have little effect on a child's beliefs. In contrast, a complier child whose parent is allowed DI because of assignment to a lenient judge may infer that applying for DI is more worthwhile since it eventually leads to success, even if the case is marginal.

Policy relevance

Despite the local nature of our estimates, the intergenerational link among the compliers to our instrument could be relevant for policy, since reforms aimed at stemming the rise in DI will likely have the largest effect on applicants on the margin of program entry. Furthermore, in both Norway and the U.S., the rise in DI rolls in recent decades appears to be primarily driven by a more liberal screening of marginal applicants who are often initially denied and relatively likely to appeal (Autor and Duggan, 2006; Kostol and Mogstad, 2014). To illustrate the policy relevance of our findings, Appendix C simulates the total reduction in DI participation from a policy which makes the screening process more stringent by making judges less likely to allow an appeal.

There are two components to the total reduction in DI from the policy change: the direct effect on parents, and the indirect effect on children. To calculate how the direct and indirect effects of a policy change would lower DI participation over time, we shift the value of our instrument, the judge leniency variable, downward by one-fifth of a standard deviation. Given the local nature of our estimates, our simulated policy effect only reflects the instrument-induced change in DI participation of complier children and their parents, and there is no change in the participation rates of always takers or never takers. As shown in Appendix C, this simulation suggests that in the early years after a tightening of the screening process, most of the reduction in DI participation can be attributed to the direct effect on parents, as there is little opportunity for children to learn and respond to their parent's DI experience. In contrast, the intergenerational effect grows over time; after ten years, the increase in children's participation accounts for almost half of the total reduction in DI rolls. In terms of program expenditure, it is important to capture this intergenerational effect, since few individuals exit DI after entering and the children are much younger than their parents when they enter DI.

5 The Breadth and Nature of Welfare Cultures in DI Receipt

We think of spillover effects in welfare receipt within families or other social networks as measures of welfare culture, with the understanding that culture may operate through information, beliefs or norms. Our rich data allows us to take several steps to explore the breadth and nature of welfare cultures in the context of the DI program.

5.1 Spillover Effects in Other Social Networks

Is the causal intergenerational link we estimated in Section 4 unique to parents and their children, or do links exist in other networks as well? To investigate this question, we use our research design to examine whether there are causal spillovers in three other social networks: close neighbors, spouses and siblings. For instance, we consider families in which an adult sibling appeals an initially denied DI claim. Using judge leniency as an instrument, we estimate the impact of this sibling being allowed DI during the appeals process on the probability that other adult siblings subsequently apply for and are awarded DI. This model of spillover effects across siblings can be represented by the two-equation system given in (4) and (5), except that A_i is now an indicator for whether the sibling is allowed DI in the appeal process while P_i is an indicator variable for whether the other sibling subsequently participates in DI. In a similar fashion, we examine spillover effects across spouses and close neighbors. As for siblings, we limit the sample to neighbors or spouses where one adult individual appeal an initially denied DI claim. Using street addresses, we define the four closest addresses on each side as neighbors.²⁰ Appendix Table A.6 provides descriptive statistics for these other networks.

As highlighted in Dahl, Loken, and Mogstad (2014), this research design allows us to address the well-known problems of reflection, correlated unobservables and endogenous network membership. The presence of an instrument, which appears in equation (4) but not (5), solves the reflection problem of simultaneity. Moreover, since z_i is orthogonal to all observed and unobserved covariates, correlated unobservables cannot bias the estimates of spillover effects. And finally, since we measure social networks before the realization of z_i , endogenous network membership does not create a bias either; any changes in network membership which happen after the allowance decision are either a causal result of changes in z_i or orthogonal to changes in z_i .

Appendix Table A.7 tests for random assignment of cases to judges in these other networks. For each network, we regress the judge leniency measure on a vector of appellant characteristics. As we found for our parent-child sample in Table 1, there is strong empirical support for the claim that judges are randomly

²⁰While not reported, defining neighbors more narrowly or more broadly yields similar results.

assigned. Of the 50 estimated coefficients across the three networks, only one is significant at the 5% level. Moreover, for each network, the explanatory variables are never jointly significant predictors of judge leniency.

Table 6 show estimates of spillover effects in the three networks. Our findings point to a special link between parents and their children, while there is little if any impact of close neighbors' DI receipt. The IV estimate for neighbors is close to zero (-.003), an estimate which is relatively precise due to the large sample of neighbors. In comparison, we have less precision to draw firm conclusions about spillovers in DI receipt across siblings or spouses.

Table 6: Estimates of Spillover Effects in Other Networks.

	<i>Second peer on DI 5 years after first peer's appeal decision</i>			
	First stage	Reduced form	IV	N
A. Close Neighbors				
Neighbor's judge leniency	0.746*** (0.089)	-0.003 (0.009)		161,569
Neighbor allowed DI			-0.003 (0.012)	
Dependent mean	0.10	0.06		
B. Spouses				
Spouse's judge leniency	0.834*** (0.089)	0.044 (0.082)		5,763
Spouse allowed DI			0.052 (0.096)	
Dependent mean	0.10	0.12		
C. Siblings				
Siblings's judge leniency	0.749*** (0.094)	-0.005 (0.039)		17,706
Sibling allowed DI			-0.006 (0.052)	
Dependent mean	0.10	0.10		

***p<.01, **p<.05, *p<.10. Standard errors (in parentheses) are clustered at the judge level.

Notes: The samples are restricted to the networks of individuals who appeal an initially denied DI claim during the period 1989-2005 (see Section 3 for further details). The samples are further restricted to individuals no older than 55 at the appeal decision. Panel A: Using street addresses the year before appeal, we define the four closest addresses on each side as close neighbors. There are 80 different judges. Panel B: Spouses defined as those married to the appellant the year before appeal. There are 76 different judges. Panel C: Siblings defined as those with the same parent as the appellant. There are 77 different judges. The specifications include additional controls as reported in panel B of Table 3. Online Appendix Table A.6 provides summary statistics for the three samples.

5.2 Types of Parent-Child Links

In Table 7, we examine how the intergenerational transmission of DI receipt depends on the type of parent-child link. In our main analysis, we restricted the sample to children who are age-eligible for DI (at least 18

years old) at the time of the parent’s appeal decision. Because few parents with young children apply for DI, the baseline sample includes the typical parent-child links (see Appendix Figure A.2). In Table 7, we find the intergenerational relationship remains strong even when we exclude children who live at home or focus on children who are at least 25 years of age. When we look at an alternative, smaller sample of children who are under 18 at the time of their parent’s appeal decision, we still find parental DI receipt substantially increases the probability that children will subsequently apply for DI 10 years later.²¹ Taken together, these findings suggest the influence of parental DI allowance does not depend strongly on the living arrangement or age of the child.

Table 7: Intergenerational Welfare Transmission by Living Arrangement and Age of Child.

Sample	Reduced form	IV	Dep. mean	N
<i>Child on DI 5 years after parent’s appeal decision</i>				
A. Child living away from home	0.077** (0.031)	0.080** (0.031)	0.03	8,395
B. Child at least 25 years of age	0.079*** (0.030)	0.075** (0.030)	0.03	6,489
<i>Child applied for DI 10 years after parent’s appeal decision</i>				
C. Child between the ages of 8-17	0.080** (0.039)	0.095** (0.043)	0.03	4,220

***p<.01, **p<.05, *p<.10. Standard errors (in parentheses) are clustered at the judge level.

Notes: Specifications mirror the baseline results with additional controls reported in panel B of Table 3. Child residency is determined based on whether a child has a different address from their parent one year prior to the parent’s appeal.

5.3 Possible Channels for Intergenerational Welfare Transmission

Information transmission and parental investments

In our setting, there is limited scope for several welfare transmission channels which might arise in other contexts. First of all, it is unlikely that children gain relevant information about how to initially apply (e.g., which documents to fill out) or appeal a denied DI case (e.g., how to write up the case) from parents who are allowed DI. This type of information transmission is unlikely to be important for our findings because both allowed and denied parents go through the same application and appeals process. Indeed, our experimental research design ensures that parents assigned to lenient versus strict judges have, on average, the same information to transmit to their children about how to apply or appeal.

²¹For panel C, since the children are young and not yet eligible for DI when their parent’s appeal decision occurs, we look at application 10 years later. Estimates using participation as the outcome yield large point estimates of intergenerational transmission, but the standard errors are too large to draw firm conclusions.

Another channel which is unlikely to explain our results is parental investment while a child is young. Since we restrict the baseline sample to parents who first apply for DI when their children are older, this rules out changes in childhood investments as the explanation for our findings. There could be changes in parental investments for adult children, but the results in Table 7 suggest this is unlikely to be a key factor. It is important to recognize that these channels could be important in broader samples or for other welfare programs, but the nature of our setting largely rules them out.

Changed attitudes about participation

One possible explanation for our findings is that observing a parent on DI could affect a child’s attitudes about stigma or the benefits of participation. Changes in attitudes could be global to any type of social assistance or specific to DI. Thinking about global changes, parental DI allowance could change attitudes and perceived stigma about the relative merits of work versus any type of government assistance. Table 8 empirically investigates whether traditional welfare use by a child change after a parent is allowed DI. As before, we use judge leniency as an instrument for parental DI allowance. For comparison, the first specification copies our baseline estimates for a child’s DI participation, which are large and statistically significant. The second specification regresses a child’s participation in Norway’s social assistance program (traditional welfare) on their parent’s DI allowance. This program is considered a last-resort safety net and there are no clear rules regarding eligibility or benefit amounts, with discretion being left to the local social worker. Appendix Figure A.6 reports survey evidence showing that participation in this program is highly stigmatized. Yet, both the reduced form and the IV estimates are small and statistically insignificant.²²

Table 8: Effect of Parent DI Allowance on Child Participation in Various Welfare Programs.

Dependent variable	<i>5 years after parent’s appeal decision</i>		
	Reduced form	IV	Dep. mean
A. Child on disability insurance	0.052** (0.020)	0.060*** (0.023)	0.03
B. Child on social assistance (i.e., traditional welfare)	0.001 (0.049)	0.001 (0.056)	0.10

***p<.01, **p<.05, *p<.10. Standard errors (in parentheses) are clustered at the judge level.

Notes: Baseline sample of 14,722 observations (see Table 3). There are 79 different judges. Specifications mirror the baseline results with additional controls reported in panel B of Table 3. Social assistance is a means-tested program for individuals with very low income.

²²The close to zero estimates are unlikely to reflect benefit substitution, as the correlation between DI and social assistance are slightly positive both in our sample (correlation = 0.07) and in the population at large (correlation = 0.10).

Changed beliefs about the likelihood of success

Another possible mechanism is that DI receipt by a parent changes a child’s belief about the efficacy of trying to apply for DI. A child whose parent draws a more lenient judge may infer that applying for DI is more worthwhile since it eventually leads to success, while the child of a strict judge may infer the process is unfair and not worth the effort. This is a transmission of beliefs about the efficacy of trying to get on to the program, rather than a change in beliefs about the stigma or value of being on DI itself. Our findings are consistent with both of these mechanisms, and we cannot completely sort out the direct effects on a child’s attitudes about the stigma or value of DI participation from beliefs about the efficacy of applying.

We can, however, look into whether children use their parent’s allowance or denial decision to update their beliefs about the relative likelihood of being allowed DI depending on the disorder they report. Let the probability a child applies for DI and reports the same disorder as the parent be denoted $Pr(\text{Apply} \cap \text{Same}) = Pr(\text{Same}|\text{Apply})Pr(\text{Apply})$. To see how the leniency of a parent’s judge (z) affects this probability, take the total derivative:

$$\underbrace{\frac{dPr(\text{Apply} \cap \text{Same})}{dz}}_{\text{net effect}} = \underbrace{Pr(\text{Same}|\text{Apply})\frac{dPr(\text{Apply})}{dz}}_{\text{applying}} + \underbrace{Pr(\text{Apply})\frac{dPr(\text{Same}|\text{Apply})}{dz}}_{\text{reporting}}. \quad (6)$$

Equation (6) highlights that the net effect of judge leniency on the probability of applying with the same disorder as the parent consists of two distinct components: the effect on applying given the likelihood of reporting the same disorder, and the effect on reporting the same disorder given the likelihood of applying.

Panel A of Table 9 reports results related to the decomposition in equation (6). For comparison, the first column copies from Table A.3 the impact of a parent’s judge leniency on the probability a child applies for DI, $dPr(\text{Apply})/dz$. The second column regresses the probability a child applies with the same disorder as their parent on the parent’s judge leniency. This implies a majority (63 percent) of children who were induced to apply because of a lenient judge also report the same disorder as their parent. The third column reports how much of the effect in column (2) is attributable to reporting versus applying (see equation 6). The reporting effect accounts for roughly half of the net effect of judge leniency on the probability of applying with the same disorder as the parent. This finding is consistent with an updating story where children whose parents are assigned a lenient judge revise their beliefs upward about the relative likelihood of being allowed DI if they report the same disorder as their parent.

In Panel B of Table 9, we perform a similar decomposition for participation instead of application. The results are qualitatively similar to those reported in Panel A. As expected, there is no increase in the role of reporting. This is consistent with the fact that there is no actual information about the benefit of reporting

the same disorder as one’s parent if the parent was allowed DI merely because they received a more lenient judge by chance. Any changes in beliefs about the relative likelihood of being allowed DI depending on the disorder they report is misguided.

Table 9: Do Children Apply with the Same Disorder as Their Parents?

5 years after parent’s appeal decision

A. DI application			
	<i>Apply</i>	<i>Apply \cap Same</i>	<i>Amount of column 2 explained by reporting</i>
Parent’s judge leniency	0.057*** (0.020)	0.036** (0.017)	53.0 %
Dependent mean	0.032	0.010	

B. DI participation			
	<i>Participation</i>	<i>Participation \cap Same</i>	<i>Amount of column 2 explained by reporting</i>
Parent’s judge leniency	0.052** (0.020)	0.024* (0.014)	49.6 %
Dependent mean	0.027	0.007	

***p<.01, **p<.05, *p<.10. Standard errors (in parentheses) are clustered at the judge level.

Notes: Baseline sample of 14,722 observations (see Table 3). There are 79 different judges. Specifications in column (1) mirror the baseline results with additional controls reported in panel B of Table 3. Applying/participating with the same disorder as the parent is based on reporting the same first letter of the ICD-10 code. The results in column (2) are the sum of the right hand side of the decomposition in equation 6. Standard errors are bootstrapped with 500 replications.

6 Conclusion

This paper provides novel evidence on intergenerational welfare transmission in a setting where we can credibly address concerns about correlated unobservables across generations. The key to our research design is that the DI system in Norway randomly assigns judges to DI applicants whose cases are initially denied. Some appeal judges are systematically more lenient, which leads to random variation in the probability an individual will be allowed DI. We utilize this exogenous variation to examine whether parents being allowed DI during the appeal process affects the probability their adult children subsequently apply for and are awarded DI. We find strong evidence that welfare receipt in one generation causes welfare participation in the next generation: when a parent is allowed DI, their adult child’s participation over the next five years increases by 6 percentage points. This effect grows over time, rising to 12 percentage points after ten years.

Our findings serve to highlight that welfare reforms can have long-lasting effects on program participation, since any original effect on the current generation could be reinforced by changing the participation behavior of their children as well. At the same time, they raises a number of questions. Is the causal link in welfare

receipt unique to parents and their children, or do links exist in other networks as well? To what extent does the intergenerational transmission of welfare receipt depend on the type of parent-child link? What is the relative importance of information, beliefs and norms in intergenerational welfare transmission? What does intergenerational welfare transmission look like in the population at large or in other settings?

Our study only scratches the surface of these important but difficult questions. We go beyond the transmission of DI receipt across generations and use our research design to examine spillovers in other social networks. Our findings point to a special link between parents and their children, with little impact of close neighbors' DI receipt. Our findings also reveal that parents' influence on childrens' decisions to apply for and take up DI is not specific to the living arrangement or age of the child. We explore how information, beliefs and norms could operate in our context and find suggestive evidence that what may change as a result of a parent being allowed DI is their children's beliefs about the efficacy of trying to get on to the DI program or their attitudes about DI participation and its stigma. However, it is important to emphasize the local nature of our findings. Evidence from other settings or populations would be useful to assess the generalizability of our findings.

References

- ABADIE, A. (2003): “Semiparametric Instrumental Variable Estimation of Treatment Response Models,” *Journal of Econometrics*, 113(2), 231–263.
- AIZER, A., AND J. J. DOYLE (2013): “Juvenile Incarceration, Human Capital and Future Crime: Evidence from Randomly-Assigned Judges,” Working Paper 19102, National Bureau of Economic Research.
- ANTEL, J. J. (1992): “The Intergenerational Transfer of Welfare Dependency: Some Statistical Evidence,” *The Review of Economics and Statistics*, 74(3), 467–473.
- ATKINSON, A., L. RAINWATER, AND T. SMEEDING (1995): “Income Distributions in OECD countries: Evidence from the Luxembourg Income Study,” *OECD Publications and Information Center*.
- AUTOR, D. (2011): “The Unsustainable Rise of the Disability Rolls in the United States: Causes, Consequences, and Policy Options,” *NBER Working Paper*, No. 17697.
- AUTOR, D., N. MAESTAS, K. MULLEN, AND A. STRAND (2011): “Does Delay Cause Decay? The Effect of Administrative Decision Time on the Labor Force Participation and Earnings of Disability Applicants,” *University of Michigan Working Paper*, No. 258.
- AUTOR, D. H., AND M. G. DUGGAN (2003): “The Rise in the Disability Rolls and the Decline in Unemployment,” *Quarterly Journal of Economics*, 118(1), 157–205.
- (2006): “The Growth in the Social Security Disability Rolls: A Fiscal Crisis Unfolding,” *Journal of Economic Perspectives*, 20(3), 71–96.
- BEAULIEU, N., J.-Y. DUCLOS, B. FORTIN, AND M. ROULEAU (2005): “Intergenerational Reliance on Social Assistance: Evidence from Canada,” *Journal of Population Economics*, 18, 539–562.
- BERTRAND, M., E. F. P. LUTTMER, AND S. MULLAINATHAN (2000): “Network Effects and Welfare Cultures,” *The Quarterly Journal of Economics*, 115(3), 1019–1055.
- BIRD, A. (2007): “Perceptions of Epigenetics,” *Nature*, 447(24), 396–398.
- BJORKLUND, A., M. LINDAHL, AND E. PLUG (2006): “The Origins of Intergenerational Associations: Lessons from Swedish Adoption Data,” *The Quarterly Journal of Economics*, 121(3), 999–1028.
- BLACK, D., K. DANIEL, AND S. SANDERS (2002): “The Impact of Economic Conditions on Participation in Disability Programs: Evidence from the Coal Boom and Bust,” *The American Economic Review*, 92(1), 27–50.
- BLACK, S., AND P. DEVEREUX (2011): “Recent Developments in Intergenerational Mobility,” chap. 16, pp. 1487–1541. *Handbook of Labor Economics*, Elsevier.
- BORGHANS, L., A. GIELEN, AND E. LUTTMER (2013): “Social Support Substitution and the Earnings Rebound: Evidence from a Regression Discontinuity in Disability Insurance Reform,” *American Economic*

- Journal: Economic Policy*, forthcoming.
- BOUND, J. (1989): “The Health and Earnings of Rejected Disability Insurance Applicants,” *American Economic Review*, 79(3), 482–503.
- BRATBERG, E. (1999): “Disability Retirement in a Welfare State,” *Scandinavian Journal of Economics*, 101(1), 97–114.
- BRATBERG, E., O. A. NILSEN, AND K. VAAGE (2013): “Assessing the Intergenerational Correlation in Disability Pension Recipiency,” Discussion paper, Norwegian School of Economics.
- BURKHAUSER, R. V., AND M. C. DALY (2012): “Social Security Disability Insurance: Time for Fundamental Change,” *Journal of Policy Analysis and Management*, 31(2), 454–461.
- CAMPOLIETI, M. (2004): “Disability Insurance Benefits and Labor Supply: Some Additional Evidence,” *Journal of Labor Economics*, 22(4), 863–889.
- CAMPOLIETI, M., AND C. RIDDELL (2012): “Disability Policy and the Labor Market: Evidence from a Natural Experiment in Canada, 1998-2006,” *Journal of Public Economics*, 96, 306–316.
- DAHL, G. B., K. V. LOKEN, AND M. MOGSTAD (2014): “Peer Effects in Program Participation,” *American Economic Review*, forthcoming.
- DOBBIE, W., AND J. SONG (2013): “Debt relief and debtor outcomes: Measuring the effects of consumer bankruptcy protection,” *Manuscript*.
- DOYLE, J. (2007): “Child Protection and Child Outcomes: Measuring the Effects of Foster Care,” *The American Economic Review*, 97(5), 1583–1610.
- (2008): “Child Protection and Adult Crime: Using Investigator Assignment to Estimate Causal Effects of Foster Care,” *Journal of Political Economy*, 116(4), 746–770.
- DUNCAN, G., M. HILL, AND S. HOFFMAN (1988): “Welfare Dependence Within and Across Generations,” *Science*, 239(4839), 467–471.
- FRENCH, E., AND J. SONG (2013): “The Effect of Disability Insurance Receipt on Labor Supply,” *American Economic Journal: Economic Policy*, forthcoming.
- FURTADO, D., AND N. THEODOROPOULOS (2012): “Immigrant Networks and the Take-Up of Disability Programs: Evidence from U.S. Census Data,” IZA Discussion Papers 6549, Institute for the Study of Labor (IZA).
- GOTTSCHALK, P. (1996): “Is the Correlation in Welfare Participation Across Generations Spurious?,” *Journal of Public Economics*, 63(1), 1 – 25.
- GRUBER, J. (2000): “Disability Insurance Benefits and Labor Supply,” *Journal of Political Economy*, 108(6), 1162–1183.
- IMBENS, G. W., AND D. B. RUBIN (1997): “Estimating Outcome Distributions for Compliers in Instrumental

- Variables Models,” *The Review of Economic Studies*, 64(4), 555–574.
- KLING, J. R. (2006): “Incarceration Length, Employment, and Earnings,” *The American Economic Review*, 96(3), 863–876.
- KOSTOL, A. R., AND M. MOGSTAD (2014): “How Financial Incentives Induce Disability Insurance Recipients to Return to Work,” *American Economic Review*, 104(2), 624–55.
- LEE, C., AND G. SOLON (2009): “Trends in Intergenerational Income Mobility,” *The Review of Economics and Statistics*, 91, 766–772.
- LEVINE, P., AND D. ZIMMERMAN (1996): “The Intergenerational Correlation in AFDC Participation: Welfare Trap or Poverty Trap?,” Discussion Paper 1100-96, Institute for Research on Poverty.
- MAESTAS, N., K. J. MULLEN, AND A. STRAND (2013): “Does Disability Insurance Receipt Discourage Work? Using Examiner Assignment to Estimate Causal Effects of SSDI Receipt,” *American Economic Review*, 103(5), 1797–1829.
- MAZUMDER, B. (2005): “Fortunate Sons: New Estimates of Intergenerational Mobility in the United States Using Social Security Earnings Data,” *The Review of Economics and Statistics*, 87(2), 235–255.
- MILLIGAN, K. S., AND D. A. WISE (2011): “Social Security and Retirement around the World: Mortality, Health, Employment and Disability Insurance Participation and Reforms,” *Harvard Kennedy School Working paper*.
- MOFFITT, R. (1992): “Incentive Effects of the U.S. Welfare System: A Review,” *Journal of Economic Literature*, 30(1), 1–61.
- MOORE, T. J. (2011): “The Employment Effects of Terminating Disability Benefits: Insights from Removing Drug Addictions as Disabling Conditions,” *Working Paper University of Maryland*.
- OASDI-TRUSTEES-REPORT (2012): “The 2012 Annual Report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Federal Disability Insurance Trust Funds,” Discussion paper.
- OECD (2010): *Sickness, Disability and Work: Breaking the Barriers. A Synthesis of Findings across OECD Countries*. OECD Publishing.
- OREOPOULOS, P., M. PAGE, AND A. H. STEVENS (2006): “Does Human Capital Transfer from Parent to Child? The Intergenerational Effects of Compulsory Schooling,” *Journal of Labor Economics*, 24, 729–760.
- PAGE, M. (2004): “New Evidence on the Intergenerational Correlation in Welfare Participation,” chap. 10, pp. 226–244. *Generational Income Mobility in North America and Europe*, Cambridge University Press.
- PAGE, M. E., AND A. H. STEVENS (2002): “Will You Miss Me When I Am Gone? The Economic Consequences of Absent Parents,” NBER Working Papers 8786, National Bureau of Economic Research, Inc.
- PARSONS, D. O. (1991): “The Health and Earnings of Rejected Disability Insurance Applicants: Comment,” *The American Economic Review*, 81(5), 1419–1426.

- PEPPER, J. (2000): “The Intergenerational Transmission of Welfare Receipt: A Nonparametric Bounds Analysis,” *The Review of Economics and Statistics*, 82, 472–288.
- REGE, M., K. TELLE, AND M. VOTRUBA (2009): “The Effect of Plant Downsizing On Disability Pension Utilization,” *Journal of the European Economic Association*, 7(4), 754–785.
- SOLON, G., M. CORCORAN, R. GORDON, AND D. LAREN (1988): “Sibling and Intergenerational Correlations in Welfare Program Participation,” *The Journal of Human Resources*, 23(3), 388–396.
- VAN SOEST, A., T. ANDREYEVA, A. KAPTEYN, AND J. P. SMITH (2011): “Self-Reported Disability and Reference Groups,” pp. 237–264. *Investigations in the Economics of Aging*, National Bureau of Economic Research.
- VON WACHTER, T., J. SONG, AND J. MANCHESTER (2011): “Trends in Employment and Earnings of Allowed and Rejected Applicants to the Social Security Disability Insurance Program,” *The American Economic Review*, 101(7), 3308–29.

Online Appendix for “Family Welfare Cultures”

By Gordon Dahl, Andreas Ravndal Kostøl, and Magne Mogstad

Sections:

A. Tables and Graphs

B. Technical Appendix

C. Policy Simulation

A Tables and Graphs

Table A.1: Characteristics of DI recipients in Norway and the U.S.

Characteristic	Norway	U.S.
	DI Recipients	SSDI Recipients
Difficult to verify disorder	59.2 %	57.3 %
Age (at decision on initial application)	52.2	49.1
Prior earnings relative to the median	71.0 %	69.9 %

Notes: The U.S. numbers are from Maestas, Mullen, and Strand (2013), and the Norwegian numbers are drawn from the sample of DI applicants during the years 2000-2003. Difficult to verify disorder includes musculoskeletal and mental diagnoses. Prior earnings are measured 3-5 years before the application/appeal.

Table A.2: Characteristics of DI Applicants and Appellants in Norway and the U.S.

Characteristic	Norway		U.S.	
	Applicants	Appellants	Applicants	Appellants
Difficult to verify disorder	60.9 %	69.7 %	58.5 %	62.2 %
Age (at decision on initial application)	51.1	47.1	47.1	46.1
Prior earnings relative to the median	66.5 %	50.4 %	60.5 %	56.3 %

Notes: The U.S. numbers are from Maestas, Mullen, and Strand (2013), and the Norwegian numbers are drawn from the sample of DI applicants during the years 2000-2003. Difficult to verify disorder includes musculoskeletal and mental diagnoses. Prior earnings are measured 3-5 years before the application/appeal.

Table A.3: Intergenerational Welfare Transmission using DI Applications as the Outcome.

	First stage	<i>Child applied for DI 5 years after parent's appeal decision</i>		<i>Child ever applied for DI after parent's appeal decision</i>	
		Reduced form	IV	Reduced form	IV
A. No additional controls					
Parent's judge leniency	0.909*** (0.112)	0.060*** (0.021)		0.112*** (0.035)	
Parent allowed DI			0.066*** (0.022)		0.123*** (0.039)
B. With additional controls					
Parent's judge leniency	0.869*** (0.108)	0.057*** (0.020)		0.104*** (0.031)	
Parent allowed DI			0.065*** (0.023)		0.120*** (0.036)
Dependent mean	0.12	0.03		0.08	

***p<.01, **p<.05, *p<.10. Standard errors (in parentheses) are clustered at the judge level.

Notes: Baseline sample (see Table 3). Specifications mirror the baseline results with additional controls reported in panel B of Table 3, except the outcome is child DI applications instead of child DI receipt.

Table A.4: Characteristics of Marginal Applicants.

Parental characteristic	First stage	$P[X = x]$	$P[X = x complier]$	$\frac{P[X=x complier]}{P[X=x]}$
Low education	1.024*** (0.124)	0.554 (0.004)	0.653 (0.034)	1.179 (0.061)
High education	0.716*** (0.133)	0.446 (0.004)	0.368 (0.034)	0.825 (0.076)
Young	0.859*** (0.136)	0.541 (0.004)	0.535 (0.038)	0.988 (0.07)
Old	0.885*** (0.147)	0.459 (0.004)	0.467 (0.032)	1.018 (0.07)
Married	0.926*** (0.114)	0.685 (0.004)	0.73 (0.027)	1.066 (0.04)
Not married	0.735*** (0.158)	0.315 (0.004)	0.267 (0.034)	0.846 (0.107)
High labor market experience	0.977*** (0.157)	0.484 (0.004)	0.544 (0.034)	1.125 (0.07)
Low labor market experience	0.784*** (0.123)	0.516 (0.004)	0.465 (0.032)	0.902 (0.063)
Difficult to verify disorders	0.923*** (0.147)	0.643 (0.004)	0.683 (0.032)	1.062 (0.05)
Other disorders	0.753*** (0.137)	0.357 (0.004)	0.309 (0.029)	0.866 (0.08)

***p<.01, **p<.05, *p<.10. Standard errors (in parentheses) are clustered at the judge level.

Notes: First stage, marginal distribution, complier distribution and relative likelihood for different subgroups. Sample is restricted to appeals in the period 1989-2005. The first stage is estimated including the controls described in Table 3. Low education is defined as having 10 or fewer years of education, young as age 50 or less, and high labor market experience as working at least 9 out of the 10 years prior to the appeal decision. The bootstrapped standard errors (in parentheses) in columns 2, 3 and 4 are obtained using 500 replications.

Table A.5: Effect of Judge Allowance on Child DI Participation: Age Re-weighted Estimates.

	<i>Years since court decision:</i>				
	2 years	4 years	6 years	8 years	10 years
A. Baseline					
Judge leniency	0.034***	0.067***	0.103***	0.113***	0.120***
in parent's case	(0.012)	(0.017)	(0.024)	(0.025)	(0.033)
Dependent mean	0.010	0.018	0.030	0.046	0.063
B. Age Re-weighted					
Judge leniency	0.031*	0.070***	0.110***	0.129***	0.124***
in parent's case	(0.016)	(0.020)	(0.030)	(0.026)	(0.031)
Dependent mean	0.010	0.017	0.029	0.044	0.060

***p<.01, **p<.05, *p<.10. Standard errors (in parentheses) are clustered at the judge level.

Notes: Balanced 10 year sample (see Appendix Table C.1). For comparison, Panel A presents estimates from the top panel of Appendix Table C.1. In Panel B these regressions are re-estimated after re-weighting individual observations so that the age distribution in each year is kept constant and centered around a mean age of 30.

Table A.6: Descriptive Statistics: Other Networks.

Characteristic	Appellant		
	A. Close Neighbor	B. Spouse	C. Sibling
Age (time of decision)	43.53	43.47	41.11
Female	0.64	0.69	0.63
Married	0.54	1	0.49
Foreign born	0.23	0.23	0.04
Less than high school	0.56	0.5	0.56
High school degree	0.35	0.39	0.35
Any college	0.09	0.1	0.09
Children below age 18 living at home	0.53	0.67	0.59
Previous earnings (\$), 1-10 years prior to decision	18,753	21,082	19,574
Years of work, 1-10 years prior to decision	6.35	6.94	6.88
Mental disorders	0.26	0.19	0.28
Musculoskeletal disorders	0.46	0.52	0.43
Circulatory system	0.02	0.03	0.02
Respiratory system	0.03	0.03	0.02
Nervous system	0.03	0.04	0.04
Endocrine diseases	0.04	0.04	0.03
DI allowed	0.1	0.1	0.1
Number of appellant observations	12,860	5,763	8,412
	Peer		
	A. Close Neighbor	B. Spouse	C. Sibling
Age (time of decision)	36.09	44.05	40.05
Female	0.49	0.31	0.48
Married	0.42	1	0.52
Foreign born	0.14	0.24	0.03
Less than high school	0.33	0.41	0.42
High school degree	0.42	0.42	0.41
Any college	0.25	0.16	0.16
Children below age 18 living at home	0.47	0.66	0.62
Previous earnings (\$), 1-5 years prior to decision	36,187	43,767	39,467
Years of work, 1-5 years prior to decision	4.24	4.03	4.31
DI recipient 5 years after decision	0.06	0.12	0.1
DI recipient any time after decision	0.14	0.25	0.21
Number of peer observations	161,569	5,763	17,706

Notes: The samples are restricted to the networks of individuals who appeal an initially denied DI claim during the period 1989-2005 (see Section 3 for further details). The samples are further restricted to individuals no older than 55 at the appeal decision. Column A: Using street addresses the year before appeal, we define the four closest addresses on each side as close neighbors. There are 80 different judges. Column B: Spouses defined as those married to the appellant the year before appeal. There are 76 different judges. Column C: Siblings defined as those with the same parent as the appellant. There are 77 different judges. Previous earnings and years of work are measured the year before appeal in the DI appellant sample and the year before decision in the DI applicant sample. Nominal values are deflated to 2005 and represented in US dollars using the average exchange rate NOK/\$ = 6. Unless otherwise stated, all parent and child characteristics are measured the year before appeal.

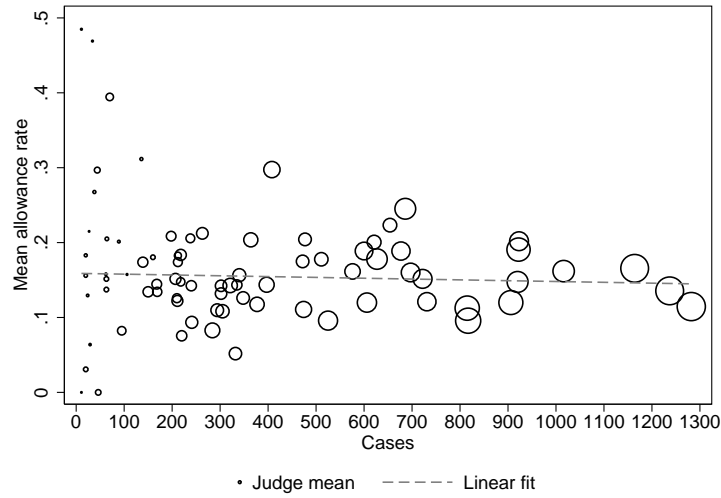
Table A.7: Testing for Random Assignment of Cases to Judges in Other Networks.

	Dependent Variable: Judge Leniency					
	A. Close Neighbors		B. Spouses		C. Siblings	
	coeff.	s.e.	coeff.	s.e.	coeff.	s.e.
Age	0.0001	(0.0001)	0.0000	(0.0001)	0.0000	(0.0002)
Female	-0.0027**	(0.0012)	0.0022	(0.0018)	-0.0003	(0.0012)
Married	0.0007	(0.0016)	—	—	-0.0001	(0.0012)
Foreign born	0.0019	(0.0018)	0.0007	(0.0023)	0.0023	(0.0040)
Less than high school	-0.0009	(0.0019)	0.0013	(0.0016)	-0.0001	(0.0014)
High school degree	0.0011	(0.0011)	-0.0008	(0.0014)	-0.0009	(0.0012)
Some college	0.0000	(0.0039)	-0.0013	(0.0031)	0.003	(0.0025)
College graduate	-0.0039	(0.0054)	-0.0001	(0.0098)	-0.0013	(0.0066)
Children below 18	0.0007	(0.0012)	0.0003	(0.0012)	-0.0007	(0.0013)
Average indexed earnings	0.0000	(0.0000)	0.0000	(0.0000)	0.0000	(0.0000)
Experience	-0.0002	(0.0002)	0.0000	(0.0002)	-0.0002	(0.0004)
Mental disorders	0.0001	(0.0036)	0.002	(0.0028)	-0.0009	(0.0028)
Musculoskeletal disorders	0.0029	(0.0036)	-0.0009	(0.0039)	0.0004	(0.0033)
Circulatory system	-0.0022	(0.0068)	-0.0055	(0.0057)	0.0031	(0.0068)
Respiratory system	0.0047	(0.0050)	0.0042	(0.0041)	0.002	(0.0037)
Nervous system	-0.0015	(0.0046)	0.0028	(0.0043)	0.0018	(0.0034)
Endocrine diseases	-0.0029	(0.0047)	-0.0011	(0.0046)	-0.0021	(0.0041)
F-statistic for joint significance	0.98		0.58		0.52	
[p-value]	[.49]		[.88]		[.93]	
Number of observations	161,569		5,763		17,706	
R-squared	.002		.002		.001	

***p<.01, **p<.05, *p<.10. Standard errors (in parentheses) are clustered at the judge level.

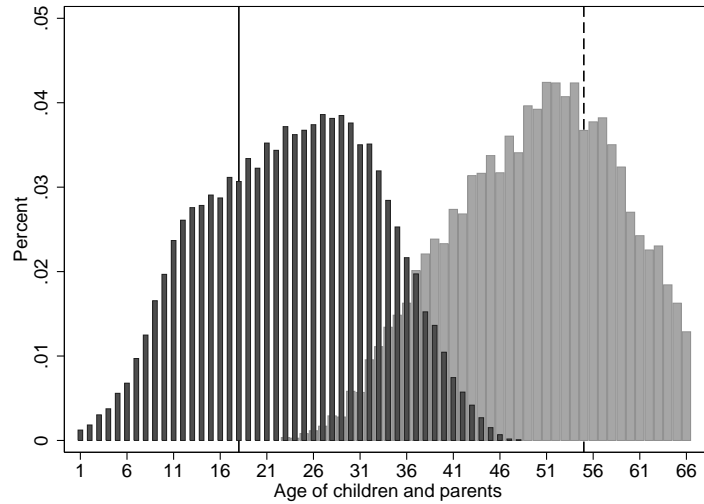
Notes: Samples are described in the notes to Appendix Table A.6. The judge leniency variable is constructed by calculating the leave-out mean judge allowance rate for all cases a judge has handled (not just those in the baseline estimation sample), regressing this measure on fully interacted year and department dummies, and using the residual from this regression as the variable. Columns 1, 3 and 5 display OLS estimates from separate regressions of judge leniency on appellant characteristics. F-statistics are obtained from OLS estimation on the combined set of applicant characteristics. All regressions include fully interacted year and department dummies. Characteristics of appellants are measured prior to the appeal. Number of children is the number under age 18, average indexed earnings is mean earnings for the last ten years prior to appeal and experience is number of years with positive earnings over this ten year period.

Figure A.1: Judge Leniency versus Number of Cases Handled.



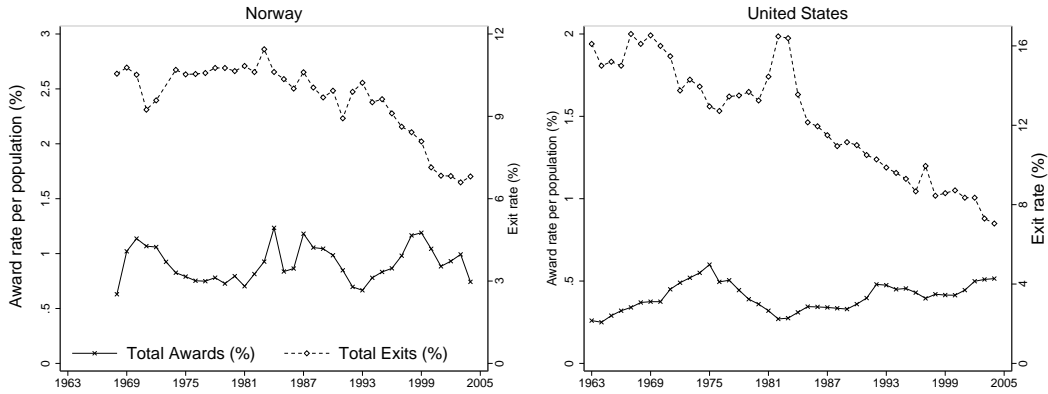
Notes: The figure plots a judge's allowance rate against the total number of cases he or she has handled. There are 79 different judges, and on average, each judge has handled a total of 380 cases. Allowance rates normalized by subtracting off year \times department deviations from the overall mean. The sample is restricted to individuals appealing their first denied case during the period 1989-2005. Dot size is proportional to the number of cases a judge handles in the estimation sample (which is weakly smaller than the number of cases they have ever handled, as plotted on the x-axis).

Figure A.2: Age Distribution of Parents and Children.



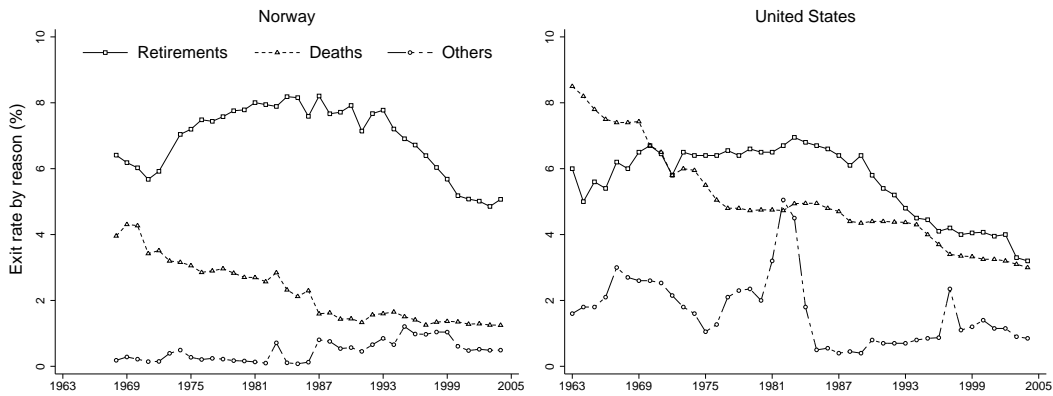
Notes: Age distribution for children (black) and parents (gray) meeting our other, non-aged based, sample restrictions. The baseline estimation sample is restricted to parents who are younger than 55 at the time of their appeal (denoted by the dashed vertical line) with children who are 18 or older (denoted by the solid vertical line).

Figure A.3:
Award and Exit Rates



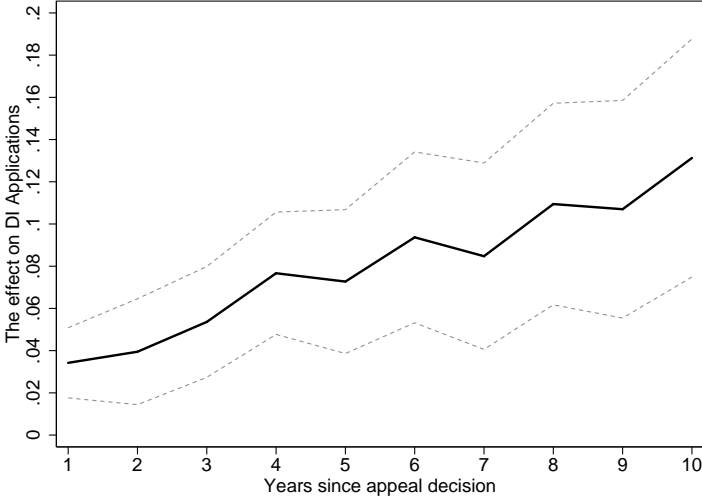
Notes: The U.S. trends are based on Autor and Duggan (2006), while the Norwegian trends are collected from various issues of the SSA Supplement. The graphs show award rates in the insured population and exit rates from the DI program in both countries.

Figure A.4:
Exit rates by reason



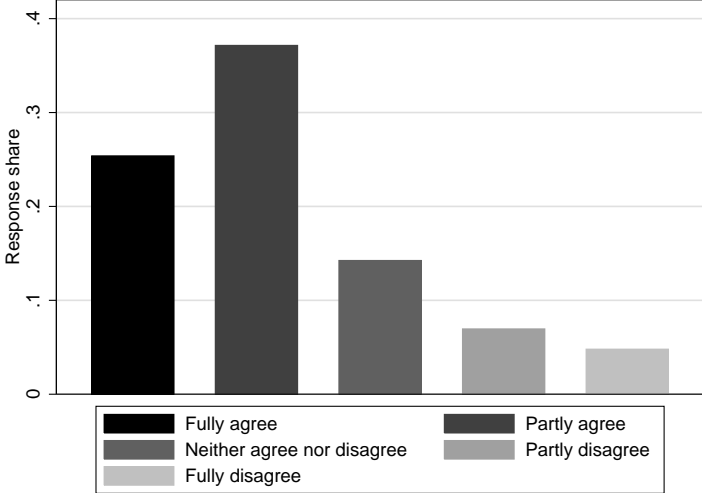
Notes: The U.S. trends are based on Autor and Duggan (2006), while the Norwegian trends are collected from various issues of the SSA Supplement. The graphs show exit rates because of death, retirement or other reasons (including eligibility-based exits).

Figure A.5: Intergenerational Transmission over Time using DI Application as the Outcome.



Note: This figure mirrors that of Figure 4 with the outcome being child DI application instead of child DI receipt. Dashed lines represent 90 percent confidence intervals (clustered at the judge level).

Figure A.6: Survey Evidence on Social Assistance and Stigma.



Notes: Responses to the statement “Receiving social assistance makes people feel like second rate citizens” from a random sample of 3,190 Norwegians in 2007. Responses of “Don’t know” (approximately 11%) are omitted from the figure.

B Technical Appendix

Let $A_i(k)$ be the potential allowance decision of child i 's parent if the value of the instrument z_i is equal to k ; this indicator variable is equal to 1 if a the parent is allowed DI and 0 otherwise. Let $P_i(A_i, z_i)$ denote child i 's potential participation in DI given parental allowance decision and the value of the instrument; this indicator variable is equal to 1 if the child participates in DI and 0 otherwise. Suppose the usual IV assumptions hold:

EXCLUSION: $P_i(A_i, z_i) = P_i(A_i)$ for $A_i = 0, 1$

INDEPENDENCE: $[P_i(1), P_i(0), \{A_i(k); \forall k\}] \perp z_i$

FIRST STAGE: $E[A_i(k) - A_i(k - 1)] \neq 0; \forall k$

MONOTONICITY: $A_i(k) \geq A_i(k - 1); \forall k$

To simplify notation, we assume that z_i is unconditionally random (alternatively, let z_i be the residual from a regression of judge leniency on fully interacted year and department dummies).

Compliance types in our sample

The goal is to calculate the shares of children whose parents are always takers, never takers and compliers in our sample. These compliance types are usually defined in the context of binary instruments, whereas judge leniency takes multiple values. In the spirit of Imbens and Rubin (1997) and Abadie (2003), we define compliers as parental appellants that could have received a different allowance decision had their case been assigned to a different judge, i.e.,

$$\pi_c \equiv Pr(A_i = 1 \mid z_i = \bar{z}) - Pr(A_i = 1 \mid z_i = \underline{z}) = Pr(A_i(\bar{z}) > A_i(\underline{z})), \quad (7)$$

where \bar{z} and \underline{z} denote the maximum (most lenient judge) and minimum (strictest judge) values of the instrument. The always takers are the parental appellants that would be allowed DI regardless of the judge assigned to their case. Because of monotonicity and independence, the share of always takers is given by the probability of allowance for the strictest judge, i.e.,

$$\pi_a \equiv Pr(A_i = 1 \mid z_i = \underline{z}) = Pr(A_i(\bar{z}) = A_i(\underline{z}) = 1). \quad (8)$$

The remaining parents are never takers; they would never be allowed DI no matter which judge that is assigned to their case, i.e.,

$$\pi_n \equiv Pr(A_i = 0 \mid z_i = \bar{z}) = Pr(A_i(\bar{z}) = A_i(\underline{z}) = 0). \quad (9)$$

In Panel A of Table B.1, we calculate the shares of children whose parents are always takers, never takers and compliers in our sample. To this end, we need an estimate of the relationship between A_i and z_i combined with values for \bar{z} and \underline{z} . In the first three columns, we use a flexible analog to the first stage equation (4). As in Figure 3, we perform a local linear regression of parental allowance on our measure of judge leniency (which controls for fully interacted year and department dummies). Using the leniency measures for the top 1 percentile (most lenient) and the bottom 1 percentile (strictest) of judge leniency, we find a 25 percent share for children of compliers, whereas children of never takers and always takers make up about 72 and 3 percent of the sample, respectively.

The last three columns shows that these calculations are quite similar if we instead use our linear specification of the first stage, given by equation (4). In this case, we can calculate π_c from $\hat{\gamma}(\bar{z} - \underline{z})$, π_a from $\hat{\alpha} + \hat{\gamma}\underline{z}$, and π_n from $1 - \hat{\alpha} - \hat{\gamma}\bar{z}$, where $\hat{\gamma}$ and $\hat{\alpha}$ are the estimated coefficients from our first stage estimates. Using the leniency measures for the top 1 percentile (most lenient) and the bottom 1 percentile (strictest) of judge leniency, we then find a complier children share of 25 percent, whereas children of never takers and always takers make up about 74 and 1 percent of the sample, respectively. By comparing the figures across the columns, it is clear that the relative size of these groups are fairly robust to the choice of model and the exact cutoffs for the most lenient and strictest judge.

Potential DI participation in our sample

Our IV estimates tell us that the probability a complier child has ever been on DI increases by 12 percentage points if their parent is allowed DI upon appeal. A natural question is: how many complier children would have been on DI if their parents had been denied DI? To answer this question, we need to calculate the sample analog of $E(P_i(0) | A_i(\bar{z}) > A_i(\underline{z}))$.

Consider children with $A_i = 0$. Those whose parents were assigned to judges with $z_i = \underline{z}$ are a mixture of children of compliers and never takers:

$$E(P_i | A_i = 0, z_i = \underline{z}) = \frac{\pi_c}{\pi_c + \pi_n} E(P_i(0) | A_i(\bar{z}) > A_i(\underline{z})) + \frac{\pi_n}{\pi_c + \pi_n} E(P_i(0) | A_i(\bar{z}) = A_i(\underline{z}) = 0)$$

while those whose parents were assigned to judges with $z_i = \bar{z}$ are children of never takers:

$$E(P_i | A_i = 0, z_i = \bar{z}) = E(P_i(0) | A_i(\bar{z}) = A_i(\underline{z}) = 0).$$

By disentangling this mixture, we can identify:

$$E(P_i(0) | A_i(\bar{z}) > A_i(\underline{z})) = \frac{\pi_c + \pi_n}{\pi_c} E(P_i | A_i = 0, z_i = \underline{z}) - \frac{\pi_n}{\pi_c} E(P_i | A_i = 0, z_i = \bar{z}). \quad (10)$$

Table B.1: Technical Appendix Calculations

Model specification:	Local linear model			Linear model		
Leniency measures, top and bottom:	1 %	1.5 %	2 %	1 %	1.5 %	2 %
A. Sample share by compliance type:						
Children of compliers	0.25	0.23	0.21	0.25	0.22	0.20
Children of never takers	0.72	0.74	0.75	0.74	0.76	0.77
Children of always takers	0.03	0.03	0.04	0.01	0.02	0.03
B. Potential participation rate (Ever on DI):						
Complier children if parents are denied	0.026	0.029	0.038	0.003	0.004	0.007
C. OLS decomposition (Ever on DI):						
Potential participation of treated children	0.081	0.081	0.083	0.080	0.080	0.084
Potential participation of untreated children	0.077	0.077	0.078	0.076	0.076	0.075
Implied OLS from potential participation rates	0.004	0.004	0.006	0.004	0.004	0.009
Actual OLS estimate (std. error)				0.007 (0.007)		

In Panel A of Table B.1, we calculated the shares of children whose parents are always takers, never takers and compliers in our sample. To compute the remaining quantities in equation (10), we need an estimate of the relationship between P_i and z_i in the subsample of children with $A_i = 0$. In the first three columns of Panel B in Table B.1, we perform a local linear regression of child DI participation on our measure of judge leniency (which controls for fully interacted year and department dummies) in this subsample. Using the leniency measures for the top 1 percentile (most lenient) and the bottom 1 percentile (strictest) of judge leniency, we calculate that 2.6 percent of the complier children would have been on DI if their parents had been denied DI. The last three columns shows the potential participation rate if we instead use a linear probability model to estimate the relationship between child DI participation and our measure of judge leniency. These linear estimates point to even fewer complier children that would have participated in DI if their parents' appeal had been denied. By comparing the figures across the columns, it is clear that the conclusion of a relatively low value of $E(P_i(0) | A_i(\bar{z}) > A_i(\underline{z}))$ in our sample is robust to the choice of model and the exact cutoffs for the most lenient and strictest judge.

Comparison with OLS

The OLS estimand is given by:

$$E[P_i | A_i = 1] - E[P_i | A_i = 0] = E[P_i(1) | A_i = 1] - E[P_i(0) | A_i = 0]$$

The treated consist of always-takers and treated compliers, whereas the untreated consists of never-takers and untreated compliers. Below, we decompose OLS into a weighted average of the potential participation

rates of these groups. We derive this decomposition for the general case where the instrument takes $K + 1$ discrete values, $z_i \in 0, 1, \dots, K$.

Treated:

The treated consists of individuals with $A_i(0) = 1$ or $\{A_i(k) > A_i(k - 1) \text{ and } z_i \geq k; \forall k\}$. For this group, the expected participation rate is:

$$\begin{aligned} E[P_i(1)|A_i = 1] &= E[P_i(1)|A_i(0) = 1]Pr(A_i = 1 | z_i = 0)/Pr(A_i = 1) \\ &+ \sum_{k=1}^K \{E[P_i(1)|A_i(k) > A_i(k - 1)] \times \\ &[(Pr(A_i = 1 | z_i = k) - Pr(A_i = 1 | z_i = k - 1))/Pr(A_i = 1)] \sum_{j=k}^K Pr(z_i = j)\} \end{aligned}$$

To help make this equation clearer, consider the special case in which the instrument takes three values, $z_i \in 0, 1, 2$. Given monotonicity, the treated consists of the following non-overlapping groups: $A_i(0) = 1; A_i(2) > A_i(1) \text{ and } z_i = 2; A_i(1) > A_i(0) \text{ and } z_i = 2; A_i(1) > A_i(0) \text{ and } z_i = 1$. In this special case,

$$\begin{aligned} E[P_i(1)|A_i = 1] &= E[P_i(1)|A_i(0) = 1]Pr(A_i(0) = 1|A_i = 1) \\ &+ E[P_i(1)|A_i(2) > A_i(1), z_i = 2]Pr(A_i(2) > A_i(1), z_i = 2|A_i = 1) \\ &+ E[P_i(1)|A_i(1) > A_i(0), z_i = 2]Pr(A_i(1) > A_i(0), z_i = 2|A_i = 1) \\ &+ E[P_i(1)|A_i(1) > A_i(0), z_i = 1]Pr(A_i(1) > A_i(0), z_i = 1|A_i = 1) \\ &= E[P_i(1)|A_i(0) = 1]Pr(A_i(0) = 1|A_i = 1) \\ &+ E[P_i(1)|A_i(2) > A_i(1)]Pr(A_i(2) > A_i(1), z_i = 2|A_i = 1) \\ &+ E[P_i(1)|A_i(1) > A_i(0)]Pr(A_i(1) > A_i(0), z_i = 2|A_i = 1) \\ &+ E[P_i(1)|A_i(1) > A_i(0)]Pr(A_i(1) > A_i(0), z_i = 1|A_i = 1) \end{aligned}$$

where the last equality follows from the independence assumption. The weights add up to one and can be recovered from data. For example,

$$\begin{aligned} Pr(A_i(2) > A_i(1), z_i = 2|A_i = 1) &= \frac{Pr(A_i = 1|A_i(2) > A_i(1), z_i = 2)Pr(A_i(2) > A_i(1), z_i = 2)}{Pr(A_i = 1)} \\ &= \frac{Pr(A_i(2) > A_i(1), z_i = 2)}{Pr(A_i = 1)} \\ &= \frac{[Pr(A_i = 1 | z_i = 2) - Pr(A_i = 1 | z_i = 1)]Pr(z_i = 2)}{Pr(A_i = 1)} \end{aligned}$$

Untreated:

The untreated consists of individuals with $A_i(K) = 0$ or $\{A_i(k) > A_i(k-1) \text{ and } z_i < k; \forall k\}$. For example, if the instrument takes three values $z_i \in 0, 1, 2$, monotonicity implies the untreated consists of the following non-overlapping groups: $A_i(2) = 0$; $A_i(2) > A_i(1)$ and $z_i = 1$; $A_i(2) > A_i(1)$ and $z_i = 0$; $A_i(1) > A_i(0)$ and $z_i = 0$. Using the same logic as above, the general case can be written as:

$$\begin{aligned}
 E[P_i(0)|A_i = 0] &= E[P_i(0)|A_i(K) = 0]Pr(A_i = 0 | z_i = K)/Pr(A_i = 0) \\
 &+ \sum_{k=1}^K \{E[P_i(0)|A_i(k) > A_i(k-1)] \times \\
 &\quad [(Pr(A_i = 1 | z_i = k) - Pr(A_i = 1 | z_i = k-1))/Pr(A_i = 0)] \sum_{j=0}^{k-1} Pr(z_i = j)\}
 \end{aligned}$$

Potential participation rates:

We know that $E(P_i | A_i = 1, z_i = 0) = E[P_i(1)|A_i(0) = 1]$ and $E(P_i | A_i = 0, z_i = K) = E[P_i(0)|A_i(K) = 0]$. To complete the OLS decomposition, we need to recover $E(P_i(1) | A_i(k) > A_i(k-1))$ and $E(P_i(0) | A_i(k) > A_i(k-1))$. To this end, it is useful to describe the compliance behavior of different individuals, that is how parental allowance decision depends on different values of the instrument. Consider the k th margin of the instrument, from $z_i = k-1$ to $z_i = k$:

		z_i	
		$k-1$	k
A_i	0	$A_i(k) > A_i(k-1)$ or $A_i(k) = 0$	$A_i(k) = 0$
	1	$A_i(k-1) = 1$	$A_i(k) > A_i(k-1)$ or $A_i(k-1) = 1$

The participation rates of children with $A_i = 1$ whose parents were assigned to judges with $z_i = k$ is given by:

$$\begin{aligned}
 E(P_i | A_i = 1, z_i = k) &= \frac{Pr(A_i(k) > A_i(k-1))}{Pr(A_i(k) > A_i(k-1)) + Pr(A_i(k-1) = 1)} E(P_i(1) | A_i(k) > A_i(k-1)) \\
 &+ \frac{Pr(A_i(k-1) = 1)}{Pr(A_i(k) > A_i(k-1)) + Pr(A_i(k-1) = 1)} E(P_i(1) | A_i(k-1) = 1)
 \end{aligned}$$

while the participation rate of those with $A_i = 1$ whose parents were assigned to judges with $z_i = k-1$ is equal to:

$$E(P_i | A_i = 1, z_i = k-1) = E(P_i(1) | A_i(k-1) = 1)$$

By disentangling this mixture, we can identify:

$$\begin{aligned}
E(P_i(1) | A_i(k) > A_i(k-1)) &= \frac{Pr(A_i(k) > A_i(k-1)) + Pr(A_i(k-1) = 1)}{Pr(A_i(k) > A_i(k-1))} E(P_i | A_i = 1, z_i = k) \\
&- \frac{Pr(A_i(k-1) = 1)}{Pr(A_i(k) > A_i(k-1))} E(P_i | A_i = 1, z_i = k-1)
\end{aligned}$$

Using a similar logic, we can also identify:

$$\begin{aligned}
E(P_i(0) | A_i(k) > A_i(k-1)) &= \frac{Pr(A_i(k) > A_i(k-1)) + Pr(A_i(k) = 0)}{Pr(A_i(k) > A_i(k-1))} E(P_i | A_i = 0, z_i = k-1) \\
&- \frac{Pr(A_i(k) = 0)}{Pr(A_i(k) > A_i(k-1))} E(P_i | A_i = 0, z_i = k)
\end{aligned}$$

Taken together, these results allow us to express the OLS estimand in terms of observable quantities only, both in terms of the weights and the potential participation rates.

Empirical calculations

The above expressions show how to obtain the OLS estimate by taking the weighted average of the potential participation rates of subpopulations with different compliance types. To compute the weights, we need an estimate of the relationship between A_i and z_i (giving $Pr(A_i = 1 | z_i = k) - Pr(A_i = 1 | z_i = k-1)$, $Pr(A_i = 0 | z_i = K)$, and $Pr(A_i = 1 | z_i = 0)$) combined with the distribution of z_i (giving $Pr(z_i = k)$). To calculate the potential participation rates, we also need an estimate of the relationships between P_i and z_i in the subsample of children with $A_i = 0$ (giving $E(P_i(0) | A_i(k) > A_i(k-1))$ and $E[P_i(0)|A_i(K) = 0]$) and in the subsample of children with $A_i = 1$ (giving $E(P_i(1) | A_i(k) > A_i(k-1))$ and $E[P_i(1)|A_i(0) = 1]$). Taken together, these quantities allows us to compare the sample analogs of $E[P_i(1)|A_i = 1]$ and $E[P_i(0)|A_i = 0]$, and compute the implied OLS estimate.

In Panel C of Table B.1, we perform these calculations. In the first three columns, we use local linear regressions of parental allowance/child participation on the our measure of judge leniency. The last three columns use linear probability models to estimate the relationship between parental allowance/child participation and judge leniency. We break the judge leniency measure into one percentage point bins; this yields 28 margins when trimming the data to exclude the top and bottom 1% of the data.

Our calculations reveal that a treated child's potential participation rate if their parent is allowed DI is similar in magnitude to a untreated child's potential participation rate if their parent is denied DI. As a result, the implied OLS estimate is close to zero, mirroring closely the OLS estimate we obtain when we regress child DI participation on parental DI allowance. Our calculations also reveal that a child's potential participation rate if their parent is allowed DI is smaller for always takers than for compliers, whereas a child's potential participation rate if the parent is denied DI is larger for never takers than for compliers.

Taken together, this pattern of heterogeneity generates an OLS estimate that is close to zero, especially since children of never takers make up most of the group of untreated (about 81 percent).

C Policy Simulation

Our results provide causal evidence that welfare use in one generation causes welfare use in the next generation. These intergenerational effects could have important implications for the evaluation of welfare reforms, as any changes will affect not only the current generation, but also have spillover effects on their children. In this online appendix, we simulate the total reduction in DI participation from a policy which makes the screening process more stringent.

Specifically, we consider a policy change which makes all judges one-fifth of a standard deviation less likely to allow an appeal (a change of 0.012 in our judge leniency variable), a policy which conceivably could be achieved by instructing judges to be stricter in their rulings. This change translates into the average judge being approximately 10 percent less likely to grant an appeal. Our simulation focuses on how this policy change would affect the participation rates of parents and children in the balanced 10 year sample. In particular, we abstract from any behavioral responses among parents at the margin of applying for DI; we further assume no changes in the screening by case examiners at the initial application step. Given the local nature of our estimates, the simulated policy effect only reflects the instrument-induced change in DI receipt of complier children and their parents, keeping fixed the DI rate of always takers, never takers, and their children.

There are two components to the total reduction in DI from the policy change: the direct effect on parents, and the indirect effect on children. To calculate the direct effect on parents, we regress parental DI participation in a given year on judge leniency, and multiply this estimated coefficient by (minus) one-fifth of a standard deviation. We perform a similar calculation for children, regressing child DI participation in a given year on their parent's judge leniency measure, and multiply this estimated coefficient by (minus) one-fifth of a standard deviation. We then calculate how much these direct and indirect effects would lower DI participation over time. Table C.1 displays the estimated coefficients of judge leniency on child and parental DI participation in every second year. The effect of judge leniency on child DI participation grows substantially over time. In contrast, the effect of judge leniency on parental DI participation shrinks over time. This is in part because some initially rejected parents re-apply and are awarded DI and in part because some parents reach early retirement age and exit DI.

Using the estimates in Table C.1, we graph the results of the policy simulation in Figure C.1. In the first year after the court decision, making judges one-fifth of a standard deviation less likely to allow an appeal reduces DI participation by 9 percent. Most of this initial reduction can be attributed to the direct effect on parents, as there is little opportunity for children to learn and respond to their parent's DI experience. Over time, however, the direct effect of tightening the appeals process shrinks; by year 10, the direct effect

of the policy change results in a 2 percent drop in DI rolls. In contrast, the indirect intergenerational effect grows over time. After ten years, the increase in children’s participation accounts for a 2.5 percent reduction in the DI rolls. Taken together, these results show that in the first years after making the DI program more stringent, almost all of the drop in participation is due to the fact that fewer parents are being allowed DI; but 10 years later, more than half of the reduction in DI is accounted for by the reduced participation of the children of the original applicants.

This simulation makes clear that failing to account for intergenerational effects will provide misleading projections of post-reform participation rates and program costs. To translate the participation patterns shown in Figure C.1 into cost terms, we calculated the net present value of the simulated policy change for parents and children over time, based on average DI benefit amounts and assuming a 3 percent annual discount rate. Making judges approximately 10 percent stricter decreases the net present value of program expenditures after 10 years by roughly 8 percent. Nearly three quarters of this cost reduction is due to fewer parents being on DI. But one fourth of the reduction is due to the fact that fewer children participate in DI as well. If one were to extrapolate past ten years by assuming that there are no further changes in DI take up among parents or children after year 10, and that parents and children stay on DI until they reach retirement at age 67, the contribution of children to total costs is even more important. More than 50 percent of the reduction in total costs is now accounted for by the reduction in children’s participation. This is due to the fact that children entering DI have many years left before retirement, while parents are older and age out of the system sooner.

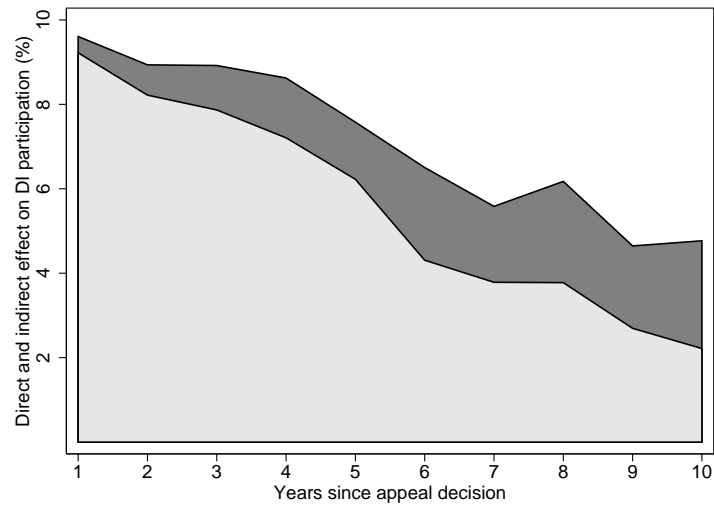
Table C.1: Effect of Judge Leniency on Child and Parent DI Participation Over Time.

	<i>Years since parent’s appeal decision</i>				
	2 years	4 years	6 years	8 years	10 years
A. Child on DI					
Judge leniency in parent’s case	0.034*** (0.012)	0.067*** (0.017)	0.103*** (0.024)	0.113*** (0.025)	0.120*** (0.033)
Dependent mean	0.01	0.018	0.03	0.046	0.063
B. Parent on DI					
Judge leniency in parent’s case	0.799*** (0.125)	0.700*** (0.138)	0.419*** (0.120)	0.367*** (0.106)	0.215** (0.082)
Dependent mean	0.411	0.562	0.661	0.726	0.773

***p<.01, **p<.05, *p<.10. Standard errors (in parentheses) are clustered at the judge level.

Notes: Balanced 10 year sample created by restricting the baseline sample to parents who appeal an initially denied DI claim during the period 1989-2000. There are 9,062 individual observations and 50 different judges. Panel A regresses child DI participation in a given year on their parent’s judge leniency. Panel B regresses parent DI participation in a given year on parent’s judge leniency. Specifications mirror column 6 of Table 3.

Figure C.1: The Effect of Tightening the Screening Process on Parents and Their Children.



Notes: Balanced 10 year sample (see Appendix Table C.1). The figure display the direct and indirect effects of tightening the screening process by one-fifth of a standard deviation. Light gray area is the direct effect on parents' participation. Dark gray area is the indirect effect on their children due to the intergenerational transmission of DI use. The estimated coefficients underlying this graph are shown for every second year in Appendix Table C.1.