FAMILY WELFARE CULTURES*

GORDON B. DAHL
ANDREAS RAVNDAL KOSTØL
MAGNE MOGSTAD

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. Our context is 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 for a causal link across generations: when a parent is allowed DI at the appeal stage, 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 10 years. Although these findings are specific to our setting, they 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. *JEL* Codes: I38, J62, H53.

I. 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 policy makers 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 that have nothing to do with a welfare culture but nonetheless translate into

*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 Brofoss, Espen Vihle, and Runar Narvland for their help in accessing the data and understanding the institutional details. The project received generous financial support from the Norwegian Research Council.

The Quarterly Journal of Economics (2014), 1711–1752. doi:10.1093/qje/qju019.
Advance Access publication on August 21, 2014.

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. Sorting out the relative importance of these explanations is central to understanding how economic conditions or government policies may shape the persistence of welfare receipt across generations.

Estimating whether welfare receipt in one generation causes welfare participation in the next has proven difficult given the likelihood of correlated unobservables across generations.[1] On top of this, it is difficult to access large data sets on welfare use that link family members together across generations. Black and Devereux (2011) 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" (p. 1531).

In this article, 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 that 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 United States, for example, outlays for DI exceed those for food stamps, traditional cash welfare, or the Earned Income Tax Credit.[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, DI rolls have steadily risen from less than 1% to

1. Researchers have documented strong intergenerational patterns for a variety of socioeconomic variables (for a summary, see Black and Devereux 2011). Levine and Zimmerman (1996) and Bjorklund, Lindahl, and Plug (2006) argue that income and family characteristics (both pre- and postbirth) can explain a large portion of such patterns, as well as intergenerational correlations in welfare use.

2. In 2011 the United States 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). In 2009, DI payments constituted 1.8% of GDP in the United States and 2.3% of GDP across the European OECD countries (OECD 2010).

over 5% of the adult population in the United States, from 1% to 7% in the United Kingdom, and from 2% to almost 10% in Norway. Many have argued that these increases are fiscally unsustainable, especially as current DI recipients are younger and have longer average life expectancies 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 use 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, an identification approach that 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 2014), the effects 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, is 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 10 years after the judge's decision. By comparison, we calculate that around 3% 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 allow us to take several steps to explore the breadth and nature of such welfare cultures. First, we 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 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 data set go through these processes. 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. Suggestive evidence for this comes from an analysis showing 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. Parental allowance may also change a child's attitudes about participation and its stigma; if it does, it appears to be DI-specific because it does not affect a child's receipt of social assistance (i.e., traditional cash welfare).

Our article 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] 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

3. See, for example, Bound (1989); Gruber (2000); von Wachter, Song, and Manchester (2011); Kostol and Mogstad (2014). 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); Autor and Duggan (2003); Campolieti (2004); Rege, Telle, and Votruba (2009). None of these studies consider the role played by intergenerational welfare transmission.

DI affects individuals on margins other than labor force participation. Our study provides some of the first causal evidence on what influences DI applications and the effects of DI receipt by parents on 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 United States and find that DI receipt substantially reduces earnings and employment of applicants. Exploiting the leniency of appeal judges in the United States, French and Song (2014) 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 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 instrumental variables (IV) estimates are local average treatment effects (LATEs), so we need to be cautious in extrapolating them 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, compared with 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 United States, 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 that makes the screening process more stringent illustrates that accounting for intergenerational effects can be important for accurate projections of

postreform 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 article proceeds as follows. Section II discusses the challenges in estimating intergenerational welfare transmission and our experimental research design. In Section III, we describe the data, provide institutional background, and compare the DI program in Norway with that of the United States. Section IV presents and interprets our main findings on intergenerational welfare transmission. Section V explores the breadth and nature of welfare cultures in DI receipt. The final section offers some concluding remarks. All appendix material can be found in the Online Appendix.

II. Identifying Intergenerational Welfare Transmission

II.A. Threats to Identification and Previous Research

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:

$$P_i^{c*} = \alpha^c + \beta^c P_i^p + \delta^c x_i^c + \varepsilon_i^c, \tag{1}$$

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:

$$P_i^{p*} = \alpha^p + \beta^p P_i^g + \delta^p x_i^p + \varepsilon_i^p, \tag{2}$$

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):

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

where $I(\cdot)$ is the indicator function. This formulation makes clear that if $corr(\varepsilon_i^p, \varepsilon_i^c | x_i^c, x_i^p) \neq 0$, 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, because there is a genetic component to health, certain physical ailments could reduce work capacity within families in ways unrelated to program participation.

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 $corr(\varepsilon_i^g, \varepsilon_i^c | x_i^c, x_i^p, x_i^g) \neq 0$, this can additionally bias the family welfare culture parameter. The potential for this type of bias is suggested by studies that document multigenerational correlations in a variety of variables such as income, poverty, education, and occupation (Oreopoulos, Page, and Stevens 2006; Lee and Solon 2009; Black and Devereux 2011). There is also evidence on multigenerational 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 ordinary least squares (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) provide evidence of modest but significant intergenerational correlation in DI receipt in Norway.[4]

4. For other studies of network effects in DI, see Furtado and Theodoropoulos (2012) and van Soest et al. (2011). For observational studies of other welfare programs, see Solon et al. (1988), Moffitt (1992), and Page (2004).

Whereas 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.[5] Pepper (2000) illustrates the difficulty in drawing credible inference with standard instruments and observational data.

II.B. Experimental Setting and Research Design

In this subsection, we begin by reviewing key facts regarding the DI program in Norway. We 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.

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

5. See Levine and Zimmerman (1996) and Gottschalk (1996).

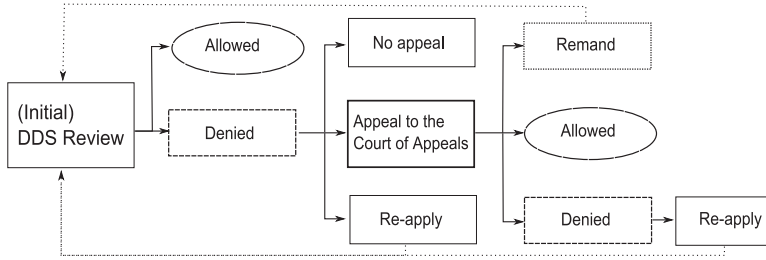

FIGURE I
DI Application and Appeals Process

The disability determination process involves multiple steps, as diagrammed in Figure I. 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 nonmedical 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 two months to the Court of Appeals. About 25% of all denials are appealed. DI appeals are reviewed by administrative law judges (ALJs). An 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 reevaluated with updated information).[6] Approximately 15% of all

6. Average processing time at the DDS stage is six months, and average processing time at the appeal stage is four months. Remands are uncommon, accounting for only 5% of appeal outcomes. In our baseline analysis, we code remanded

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 reapplying to the DDS review stage. Seventy-five percent of denied appellants eventually reapply, with 65% 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.

2. 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. Afterward, 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 ALJ 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.[7]

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 only written evidence on the appellant's case, as 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

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; the results are similar.

7. We 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). We have also presented our paper at internal seminars with the current set of judges and department heads to make sure that we have understood how the cases are handled and assessed.

based on all the cases a judge has ever handled, and not just those cases appearing in our estimation sample. On average, the 79 judges in our sample 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.

3. Verifying Random Assignment. Table I 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, conditional on fully interacted year and department dummies. The first column documents that demographic, work, and health variables are highly predictive of whether an appealed case will be allowed. The third column 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 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.[8] In fact, the point estimates are close to 0, and taken together, the variables explain only 0.35% 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 that is not true in many other countries.

A natural question is why some judges are more lenient than others. Although we do not have detailed characteristics of the judges, we do know the number of cases they have handled. Whereas experienced judges appear to be slightly less lenient,

8. The coefficient on age, while close to 0, 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.

TABLE I

TESTING FOR RANDOM ASSIGNMENT OF CASES TO JUDGES

	Dependent variable			
	Case allowed		Judge leniency	
	Coeff.	Std. err.	Coeff.	Std. err.
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 ($1,000)	0.00064***	(0.00020)	0.00000	(0.00006)
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	

Notes. Baseline estimation sample, consisting of parents who appeal an initially denied DI claim during the period 1989–2005 (see Section III 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. First and third columns 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 10 years prior to appeal, and experience is number of years with positive earnings over this 10-year period. ***$p < .01$, **$p < .05$, *$p < .10$. Standard errors (in parentheses) are clustered at the judge level.

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 are more lenient than others.

4. 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, adjusted for year and department effects, as we did for Table I.[9] As we document next, 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 two-stage least squares (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 after being denied, 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 LATEs 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.

9. Although the instrument is preestimated, 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.

As discussed in greater detail later, this suggests 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.

III. Data and Background

III.A. 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), socioeconomic data (such as years of education and earnings), and geographical identifiers. The data contain 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.[10] Following Maestas,

10. Some 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.

Mullen, and Strand (2013) and French and Song (2014), our baseline estimation excludes observations for which the assigned appeal judge has handled few cases (fewer than 10 during the period 1989–2011). 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 made during the period 1989–2005. This sample restriction allows us to observe the behavior of children for at least five years after the appeal 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 sizable 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 composed 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 V, we 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 II 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% 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% versus 7%).[11] Every child is observed for at least 5 years, some children are observed for up to 21 years; on average, a child is observed for 11 years.

III.B. Institutional Background

There are a number of similarities and a few key differences between the DI systems in the United States 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 United States than in Norway. Figure II 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 United States, the time trends are quite similar.[12] From 1961 to 2012, the rate of receipt increased from 2.2% to 9.7% in Norway and from 0.8% to 5.4% in the United States. While Norway's rate has leveled off at about 10% in recent years, the U.S. DI rate continues to rise and is projected to exceed 7% by 2018 (Burkhauser and Daly 2012).

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.[13] 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

11. By way of comparison, the rate of DI receipt is equal to 3% 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 II (except for type of disability, which is not available in both data sets).

12. The cross-country difference in DI coverage is unlikely to explain the entire discrepancy in the incidence of DI: although virtually all nonelderly adults are covered in Norway, more than 80% of all nonelderly adults are covered in the United States. 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.

13. See Autor and Duggan (2006) for a discussion of this phenomenon. In the United States, 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.

long periods. As a result, the DI exit rates in both countries have decreased in the past 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.[14] The differences in characteristics are, however, less pronounced than one might expect. For instance, almost 60% of DI recipients suffer from difficult-to-verify disorders (mental illness and musculoskeletal disorders) in both the United States and Norway (see Appendix Table A.1).

Another difference is that the appeal process plays a more important role in the United States than in Norway. While 48% of the initially rejected applicants appeal in the U.S. (French and Song 2014), only 25% of the initially rejected appeal in Norway. Success rates for appeals are much higher in the United States. Appendix Table A.2 compares the characteristics of individuals who apply for DI and those who appeal an initially denied DI claim in the two countries. In both countries, 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.

IV. Evidence on Intergenerational Welfare Transmission

IV.A. Graphical Evidence

We begin our presentation of results by providing a graphical representation of the IV approach in Figure III. In the

14. For a typical DI recipient in Norway, Kostol and Mogstad (2014) calculate the replacement rate would be 31% according to U.S. program rules and 58% according to Norwegian program rules. Factoring in health insurance coverage increases the effective replacement rate to over 50% in the United States. In Norway, all citizens are eligible for health insurance through the Social Insurance system. These calculations abstract from differences before an individual goes on to DI. In Norway, most individuals must first participate in the sick leave program, which lasts at most one year and precludes full-time work. This is an departure from the United States, where eligibility requires an impairment that is expected to last at least a year but does not require participation in any programs before DI receipt.

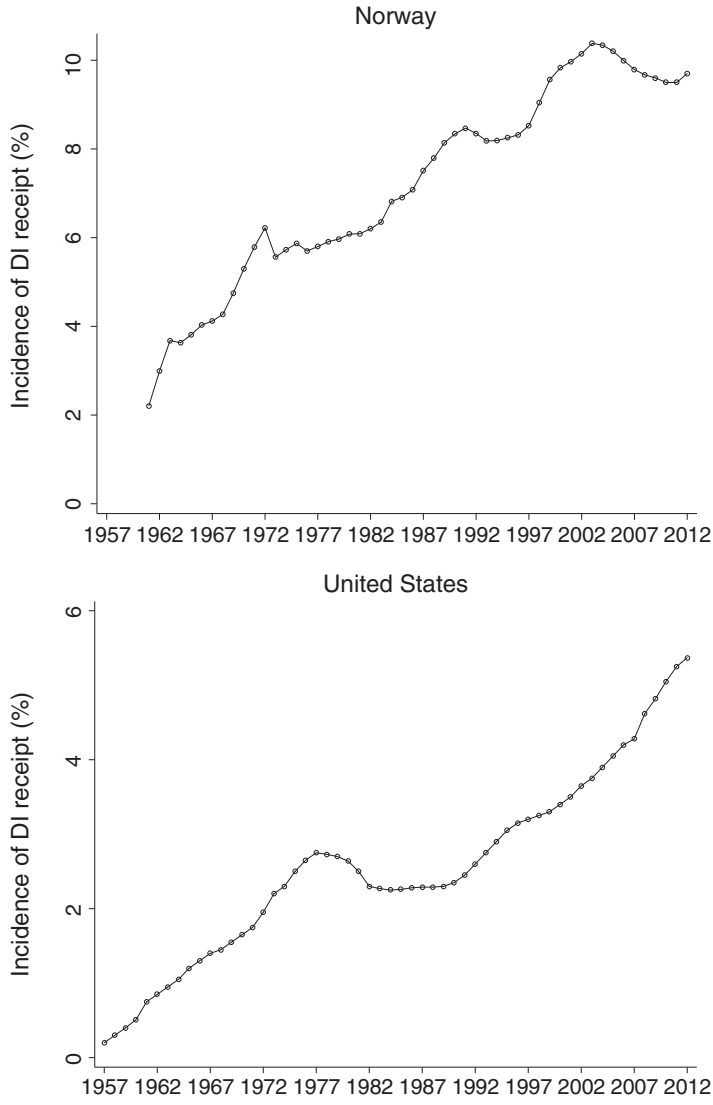

FIGURE II

Trends in DI Receipt in Norway and the United States

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 United States).

TABLE II

DESCRIPTIVE STATISTICS

	DI applicants		DI appellants	
Characteristic	Mean	Std. dev.	Mean	Std. dev.
Panel 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 ($), 1–10 years prior to decision	29,721	[22,052]	20,681	[19,037]
Years of work, 1–10 years prior to decision	7.94	[3.06]	6.78	[3.61]
Mental disorders	0.24	[0.43]	0.21	[0.41]
Musculoskeletal disorders	0.41	[0.49]	0.52	[0.50]
Circulatory system	0.06	[0.25]	0.04	[0.19]
Respiratory system	0.03	[0.16]	0.03	[0.17]
Nervous system	0.06	[0.23]	0.03	[0.17]
Endocrine diseases	0.02	[0.13]	0.04	[0.19]
DI allowed	0.75	[0.43]	0.12	[0.32]
DI recipient any time after decision	0.91	[0.28]	0.81	[0.39]
Number of parents	98,206		7,331	
Panel B: Children				
Age (time of decision)	25.33	[4.63]	24.98	[4.63]
Female	0.43	[0.49]	0.49	[0.50]
Married	0.14	[0.35]	0.16	[0.37]
Foreign born	0.17	[0.38]	0.13	[0.33]
Less than high school	0.49	[0.50]	0.52	[0.50]
High school degree	0.37	[0.48]	0.37	[0.48]
Any college	0.14	[0.35]	0.12	[0.32]
Children below age 18 living at home	0.40	[0.49]	0.31	[0.46]
Previous earnings ($), 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 older 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 U.S. dollars using the average exchange rate NOK/$ = 6. Unless otherwise stated, all parent and child characteristics are measured the year before parental application/appeal. DI receipt five years (any time) after decision refers to child DI receipt five years (any time) after the decision on the parent's case at the application (first column) or appeal (second column) stage.

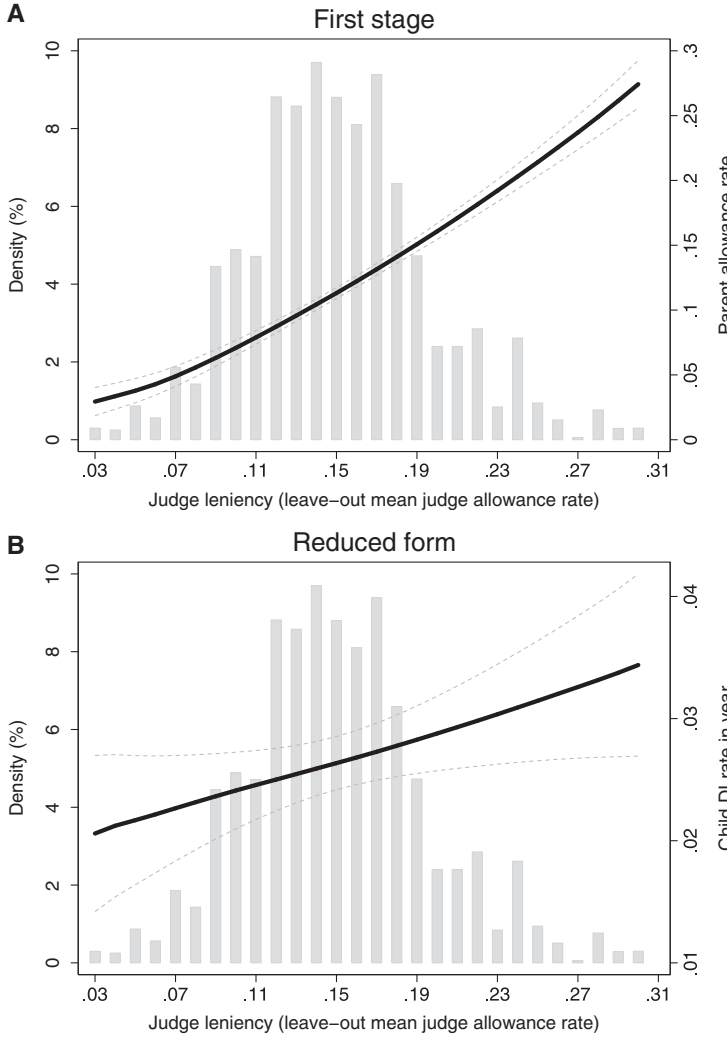

FIGURE III

Effect of Judge Leniency on Parents (First Stage) and Children (Reduced Form)

Baseline sample, consisting of parents who appeal an initially denied DI claim during the period 1989–2005 (see Section III for further details). There are 14,722 individual observations and 79 different judges. Judge leniency based on all cases a judge has ever handled, and not just the cases in our estimation sample. 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% confidence intervals.

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. The mean of the leniency variable is 0.15 with a standard deviation of 0.06. 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.

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 2.5% of children whose parents had a relatively strict judge (leniency measure = 0.09, the 10th percentile) are predicted to participate in DI five years later. This can be contrasted with roughly 3% of children whose parents had a relatively lenient judge (leniency measure = 0.22, the 90th percentile).

IV.B. Regression Estimates

We turn to a regression based analysis. The first column in Table III 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 point. This effect is not statistically different from 1.

Columns (2) and (3) of Panel A report results for whether the child participates in DI within five years after the parent's appeal decision. The second column reports the reduced-form estimate of a parent's judge leniency measure for this child outcome. The estimate of 0.055 implies that when judge leniency for a parent

TABLE III
ESTIMATES OF INTERGENERATIONAL WELFARE TRANSMISSION

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

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 III 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 10 years prior to appeal, and experience is number of years with positive earnings over this ten year period. ***$p < .01$, **$p < .05$, *$p < .10$. Standard errors (in parentheses) are clustered at the judge level.

rises by 10 percentage points, a child's DI participation will rise by roughly 0.5 percentage point. This is a sizable effect compared to the 3% average DI participation rate within five years for this sample. The third column takes the reduced-form estimate of the second column and divides it by the first-stage estimate in the first column. Since the first stage is close to 1, the reduced-form and the IV estimates are very similar.

The next two columns of the table perform a similar exercise, but look at whether the child has ever been on DI after the parent's appeal decision. Although every child is observed for at least five years after their parent's appeal decision, in these columns some children will be observed for up to 21 years; on average, the children are observed for 11 years. The unbalanced nature of columns (4) and (5) affects the interpretation of the estimates, but it should not affect their validity given the nature of our

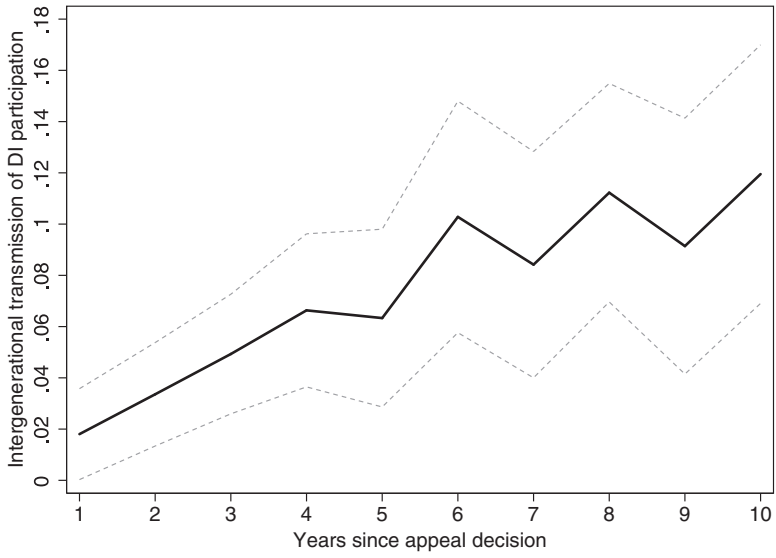

FIGURE IV

Estimates of Intergenerational Transmission over Time

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 the third column of Panel A in Table III. Dashed lines represent 90% confidence intervals (clustered at the judge level).

instrument. Figure IV complements Table III by showing IV estimates for the intergenerational transmission of DI receipt over time for a balanced panel. The estimates correspond to those in Table III, except the graph restricts the sample to children observed for at least 10 years after their parent's appeal decision.[15] Table III and Figure IV suggest the long-run effects of a parent getting on to DI are roughly twice as large as the short-run effects. For example, 10 years after the court decision, Figure IV 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

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

merely changing the timing of when their children participate in DI, at least over the time period covered by our data.

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 reestimate 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.

Last, we shift attention to how a parent's DI receipt affects the probability that their children apply for DI (as opposed to participate in DI). Appendix Figure A.5 shows IV estimates for child DI application over time based on the 10-year balanced panel used in IV. 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 13 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 article.

IV.C. Internal Validity

For judge leniency to be a valid instrument, appellants' assignment to judges must be uncorrelated with case characteristics. Table I 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 III 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 predetermined 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.[16]

Although 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, 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 only, 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 et al. 2011) and 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 he 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 IV shows that the IV estimates do not change appreciably if we control for a judge's average processing time (an exogenous variable because 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

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

TABLE IV

SPECIFICATION CHECKS FOR INTERGENERATIONAL WELFARE TRANSMISSION ESTIMATES

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

Notes. Specifications mirror the baseline results with additional controls reported in panel B of Table III. ***$p < .01$, **$p < .05$, *$p < .10$. Standard errors (in parentheses) are clustered at the judge level.

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 sizable, in line with the monotonicity assumption.

Last, Table IV 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. Although 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 fewer 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 fewer than 50 cases. Specification H considers an alternative handling of remanded cases. In our baseline analysis, we code a remanded case as rejected (see note 6). If we instead code remanded cases as allowed or denied based on eventual outcome after they are 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.

IV.D. Interpreting the IV Estimates

It is important to emphasize the local nature of our results. Our IV estimates represent a 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.

1. 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. Online 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% 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% 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% of complier children have parents with low education, while their fraction in the entire sample is only 56%. Parental compliers are also more likely to have difficult-to-verify disorders as compared to other types of parents who appeal.

2. 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 Online 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 where not allowed with the most lenient or strictest judges. We find that roughly 3% of the complier children would have ever been on DI after their parent's appeal if their parent had been denied DI, a fraction lower than the 8% observed for all children of appellants in our sample.

3. Labor and Educational Outcomes of Children. In Table V, 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 V shows that a parent's DI receipt causes employment to drop by 14 percentage points. Though we do not estimate the drop in full-time work or college completion with the same precision, both estimates suggest a sizable drop in these child outcomes as well. 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 V

EFFECT OF PARENT'S DI ALLOWANCE ON CHILD LABOR, EDUCATIONAL, AND SOCIAL ASSISTANCE OUTCOMES

	5 years after parent's appeal decision		
Dependent variable	Reduced form	IV	Dep. mean
A. DI	0.052**	0.060***	0.03
	(0.020)	(0.023)	
B. Any employment	−0.119**	−0.137**	0.58
	(0.055)	(0.065)	
C. Full-time work	−0.065	−0.075	0.42
	(0.079)	(0.090)	
D. College degree	−0.079	−0.091	0.25
	(0.060)	(0.069)	
E. Social assistance (traditional welfare)	0.001	0.001	0.10
	(0.049)	(0.056)	

Notes. Baseline sample of 14,722 observations (see Table III). There are 79 different judges. Specifications mirror the baseline results with additional controls reported in Panel B of Table III. 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. Social assistance is a means-tested program for individuals with very low income. ***$p < .01$, **$p < .05$, *$p < .10$. Standard errors (in parentheses) are clustered at the judge level.

4. 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.[17]

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 III, the OLS estimate of parental allowance on whether a child is ever on DI is

17. For example, see the discussions in French and Song (2014) 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.

0.01 (std. err. = 0.01), a number that is close to 0 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, we would expect upward-biased OLS estimates; on the other hand, OLS estimates would be biased downward 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 testing for heterogeneity in welfare transmission, exploiting that constant-effects models with a multivalued instrument and a binary endogenous regressor are overidentified. Specifically, we construct a set of dummy instruments for 10 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 reject the null hypothesis of homogeneous 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 into the potential participation rates for children of always takers, never takers, and (treated and untreated) compliers. We refer the reader to Online Appendix B for details. Our calculations reveal that a child's potential participation if their parent is allowed DI is smaller if parents are always takers than if they are compliers, whereas a child's potential participation if the parent is denied DI is larger if parents are never takers than if they are compliers. Taken together, this pattern of heterogeneity generates an OLS estimate that is close to 0, especially since children of never takers make up most of the untreated group (about 81%).

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.

5. 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 to limit the increase in DI will likely have the biggest effect on applicants on the margin of program entry. Furthermore, the recent rise in DI rolls appears to be primarily due to 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, Online Appendix C simulates the total reduction in DI participation from a policy that 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 Online 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 10 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, because few individuals exit DI after entering and the children are much younger than their parents when they enter DI.

V. 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 allow us to take several steps to explore the breadth and nature of welfare cultures in the context of the DI program.

V.A. Spillover Effects in Other Social Networks

Is the causal intergenerational link we estimated in Section IV 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 effect 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 (4) and (5), except that A_i is now an indicator for whether the sibling is allowed DI in the appeal process and 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 appeals an initially denied DI claim. Using street addresses, we define the four closest addresses on each side as neighbors.[18] 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 equation (5), solves the reflection problem of simultaneity. Moreover, because z_i is orthogonal to all observed and unobserved covariates, correlated unobservables cannot bias the estimates of spillover effects. 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 that 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 I, there is strong empirical support for the claim that judges are randomly 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 VI show estimates of spillover effects in the three networks. Our findings point to a special link between parents and their children, whereas there is little if any impact of close neighbors' DI receipt. The IV estimate for neighbors is close to 0, an estimate that 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. When interpreting these results, it is important to recognize the mechanisms at play in the various networks could differ; for example, sharing of DI benefits, joint

18. Although not reported, defining neighbors more narrowly or more broadly yields similar results.

TABLE VI
ESTIMATES OF SPILLOVER EFFECTS IN OTHER NETWORKS

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

Notes. The samples are restricted to the networks of individuals who appeal an initially denied DI claim during the period 1989–2005 (see Section III 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 III. Online Appendix Table A.6 provides summary statistics for the three samples. ***$p < .01$, **$p < .05$, *$p < .10$. Standard errors (in parentheses) are clustered at the judge level.

labor supply decisions, and leisure complementarity could be more important for spillover effects among spouses.

V.B. Types of Parent–Child Links

In Table VII, 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 VII, we

TABLE VII

INTERGENERATIONAL WELFARE TRANSMISSION BY LIVING ARRANGEMENT AND AGE OF CHILD

Sample	Reduced form	IV	Dep. mean	*N*
Child on DI 5 years after parent's appeal decision				
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
Child applied for DI 10 years after parent's appeal decision				
C. Child between the ages of 8–17	0.080** (0.039)	0.095** (0.043)	0.03	4,220

Notes. Specifications mirror the baseline results with additional controls reported in Panel B of Table III. Child residency is determined based on whether a child has a different address from their parent one year prior to the parent's appeal. ***$p < .01$, **$p < .05$, *$p < .10$. Standard errors (in parentheses) are clustered at the judge level.

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.[19] Together, these findings suggest the influence of parental DI allowance does not depend strongly on the living arrangement or age of the child.

V.C. Possible Channels for Intergenerational Welfare Transmission

1. Information Transmission and Parental Investments. In our setting, there is limited scope for several welfare transmission channels that might arise in other contexts. First, it is unlikely that children gain relevant information about how to initially apply (e.g., which documents to fill out) or appeal a

19. For 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.

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.

Another channel that is unlikely to explain our results is parental investment while a child is young. Because 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 VII 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.

2. 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. Specification E in Table V empirically investigates whether traditional welfare use by a child changes after a parent is allowed DI. As before, we use judge leniency as an instrument for parental DI allowance and regress 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.[20]

20. The close to 0 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).

3. Changed Beliefs about the Likelihood of Success. Another possible mechanism is that DI receipt at the appeal stage 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 leads to success, whereas the child of a strict judge may infer the process is unfair and not worth the effort. In particular, children of denied parents are likely to observe their parent reapply for DI, a process that is both time-consuming and costly (since to demonstrate eligibility, most parents have little earnings during the reapplication process). 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:

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

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 times the likelihood of reporting the same disorder, and the effect on reporting the same disorder conditional on applying given the likelihood of applying.

Panel A of Table VIII reports results related to the decomposition in equation (6). For comparison, the first column copies

TABLE VIII
DO CHILDREN APPLY WITH THE SAME DISORDER AS THEIR PARENTS?

	5 years after parent's appeal decision		
Panel A: DI application			
	Apply	Apply ∩ Same	% of column 2 explained by reporting
Parent's judge leniency	0.057***	0.036**	53.0
	(0.020)	(0.017)	
Dependent mean	0.032	0.010	
Panel B: DI participation			
	Participation	Participation ∩ Same	% of column 2 explained by reporting
Parent's judge leniency	0.052**	0.024*	49.6
	(0.020)	(0.014)	
Dependent mean	0.027	0.007	

Notes. Baseline sample of 14,722 observations (see Table III). There are 79 different judges. Specifications in the first column mirror the baseline results with additional controls reported in Panel B of Table III. 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 the second column are the sum of the right-hand side of the decomposition in equation (6). Standard errors are bootstrapped with 500 replications. ***$p < .01$, **$p < .05$, *$p < .10$. Standard errors (in parentheses) are clustered at the judge level.

from Appendix Table A.3 the impact of a parent's judge leniency on the probability a child applies for DI, $\frac{dPr(\text{Apply})}{dz}$. 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%) 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 the second column is attributable to reporting versus applying (see equation 6).[21] 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

21. As might be expected given how the percentages in column (2) are calculated, they are imprecisely estimated, and should therefore only be viewed as representative for our sample but not necessarily for the underlying population of appellants.

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 VIII, 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.

VI. CONCLUSION

This article 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 use 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 10 years.

Our findings highlight that welfare reforms can have long-lasting effects on program participation, because 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 raise 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 effect 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 child'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.

UNIVERSITY OF CALIFORNIA, SAN DIEGO; STATISTICS NORWAY; NATIONAL BUREAU OF ECONOMIC RESEARCH
STATISTICS NORWAY; UNIVERSITY OF BERGEN
UNIVERSITY OF CHICAGO; STATISTICS NORWAY; NATIONAL BUREAU OF ECONOMIC RESEARCH

SUPPLEMENTARY MATERIAL

An Online Appendix for this article can be found at QJE online (qje.oxfordjournal.org).

REFERENCES

Abadie, A., "Semiparametric Instrumental Variable Estimation of Treatment Response Models," *Journal of Econometrics*, 113 (2003), 231–263.
Aizer, A., and J. J. Doyle, "Juvenile Incarceration, Human Capital and Future Crime: Evidence from Randomly-Assigned Judges," NBER Working Paper 19102, 2013.
Atkinson, A., L. Rainwater, and T. Smeeding, "Income Distributions in OECD Countries: Evidence from the Luxembourg Income Study," OECD Publications and Information Center, 1995.
Autor, D., "The Unsustainable Rise of the Disability Rolls in the United States: Causes, Consequences, and Policy Options," NBER Working Paper No. 17697, 2911.
Autor, D., N. Maestas, K. Mullen, and A. Strand, "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, 2011.

Autor, D. H., and M. G. Duggan, "The Rise in the Disability Rolls and the Decline in Unemployment," *Quarterly Journal of Economics*, 118 (2003), 157–205.
———, "The Growth in the Social Security Disability Rolls: A Fiscal Crisis Unfolding," *Journal of Economic Perspectives*, 20, no. 3 (2006), 71–96.
Bertrand, M., E. F. P. Luttmer, and S. Mullainathan, "Network Effects and Welfare Cultures," *Quarterly Journal of Economics*, 115 (2000), 1019–1055.
Bird, A., "Perceptions of Epigenetics," *Nature*, 447, no. 24 (2007), 396–398.
Bjorklund, A., M. Lindahl, and E. Plug, "The Origins of Intergenerational Associations: Lessons from Swedish Adoption Data," *Quarterly Journal of Economics*, 121 (2006), 999–1028.
Black, D., K. Daniel, and S. Sanders, "The Impact of Economic Conditions on Participation in Disability Programs: Evidence from the Coal Boom and Bust," *American Economic Review*, 92 (2002), 27–50.
Black, S., and P. Devereux, "Recent Developments in Intergenerational Mobility," in *Handbook of Labor Economics* (Amsterdam: Elsevier, 2011), chapter 16, pp. 1487–1541.
Bound, J., "The Health and Earnings of Rejected Disability Insurance Applicants," *American Economic Review*, 79 (1989), 482–503.
Bratberg, E., O. A. Nilsen, and K. Vaage, "Assessing the Intergenerational Correlation in Disability Pension Recipiency," Technical report, Norwegian School of Economics, 2013.
Burkhauser, R. V., and M. C. Daly, "Social Security Disability Insurance: Time for Fundamental Change," *Journal of Policy Analysis and Management*, 31 (2012), 454–461.
Campolieti, M., "Disability Insurance Benefits and Labor Supply: Some Additional Evidence," *Journal of Labor Economics*, 22 (2004), 863–889.
Dahl, G. B., K. V. Loken, and M. Mogstad, "Peer Effects in Program Participation," *American Economic Review*, 104 (2014), 2049–2074.
Dobbie, W., and J. Song, "Debt Relief and Debtor Outcomes: Measuring the Effects of Consumer Bankruptcy Protection," Unpublished manuscript, 2013.
Doyle, J., "Child Protection and Child Outcomes: Measuring the Effects of Foster Care," *American Economic Review*, 97 (2007), 1583–1610.
———, "Child Protection and Adult Crime: Using Investigator Assignment to Estimate Causal Effects of Foster Care," *Journal of Political Economy*, 116 (2008), 746–770.
French, E., and J. Song, "The Effect of Disability Insurance Receipt on Labor Supply," *American Economic Journal: Economic Policy*, 6 (2014), 291–337.
Furtado, D., and N. Theodoropoulos, "Immigrant Networks and the Take-up of Disability Programs: Evidence from U.S. Census Data," IZA Discussion Paper 6549, 2012.
Gottschalk, P., "Is the Correlation in Welfare Participation across Generations Spurious?," *Journal of Public Economics*, 63 (1996), 1–25.
Gruber, J., "Disability Insurance Benefits and Labor Supply," *Journal of Political Economy*, 108 (2000), 1162–1183.
Imbens, G. W., and J. D. Angrist, "Identification and Estimation of Local Average Treatment Effects," *Econometrica*, 62 (1994), 467–475.
Imbens, G. W., and D. B. Rubin, "Estimating Outcome Distributions for Compliers in Instrumental Variables Models," *Review of Economic Studies*, 64 (1997), 555–574.
Kling, J. R., "Incarceration Length, Employment, and Earnings," *American Economic Review*, 96 (2006), 863–876.
Kostol, A. R., and M. Mogstad, "How Financial Incentives Induce Disability Insurance Recipients to Return to Work," *American Economic Review*, 104 (2014), 624–655.
Lee, C., and G. Solon, "Trends in Intergenerational Income Mobility," *Review of Economics and Statistics*, 91 (2009), 766–772.
Levine, P., and D. Zimmerman, "The Intergenerational Correlation in AFDC Participation: Welfare Trap or Poverty Trap?," Discussion Paper 1100-96, Institute for Research on Poverty, 1996.

Maestas, N., K. J. Mullen, and A. Strand, "Does Disability Insurance Receipt Discourage Work? Using Examiner Assignment to Estimate Causal Effects of SSDI Receipt," *American Economic Review*, 103 (2013), 1797–1829.
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., "Incentive Effects of the U.S. Welfare System: A Review," *Journal of Economic Literature*, 30 (1992), 1–61.
OASDI Trustees Report, "The 2012 Annual Report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Federal Disability Insurance Trust Funds," Technical report, 2012.
OECD, *Sickness, Disability and Work: Breaking the Barriers. A Synthesis of Findings across OECD Countries* (Paris: OECD, 2010).
Oreopoulos, P., M. Page, and A. H. Stevens, "Does Human Capital Transfer from Parent to Child? The Intergenerational Effects of Compulsory Schooling," *Journal of Labor Economics*, 24 (2006), 729–760.
Page, M., "New Evidence on the Intergenerational Correlation in Welfare Participation," in *Generational Income Mobility in North America and Europe* (Cambridge: Cambridge University Press, 2004), pp. 226–244.
Pepper, J., "The Intergenerational Transmission of Welfare Receipt: A Nonparametric Bounds Analysis," *Review of Economics and Statistics*, 82 (2000), 472–288.
Rege, M., K. Telle, and M. Votruba, "The Effect of Plant Downsizing on Disability Pension Utilization," *Journal of the European Economic Association*, 7 (2009), 754–785.
Solon, G., M. Corcoran, R. Gordon, and D. Laren, "Sibling and Intergenerational Correlations in Welfare Program Participation," *Journal of Human Resources*, 23 (1988), 388–396.
van Soest, A., T. Andreyeva, A. Kapteyn, and J. P. Smith, "Self-reported Disability and Reference Groups," in *Investigations in the Economics of Aging* (Cambridge, MA: National Bureau of Economic Research, 2011), pp. 237–264.
von Wachter, T., J. Song, and J. Manchester, "Trends in Employment and Earnings of Allowed and Rejected Applicants to the Social Security Disability Insurance Program," *American Economic Review*, 101 (2011), 3308–3329.