

Peer Effects in Program Participation

Gordon B. Dahl*

Katrine V. Løken[†]

Magne Mogstad[‡]

June 22, 2012

Abstract: The influence of peers could play an important role in the take up of social programs. However, estimating peer effects has proven challenging given the problems of reflection, correlated unobservables, and endogenous group membership. We overcome these identification issues in the context of paid paternity leave in Norway using a regression discontinuity design. Our approach differs from existing literature which attempts to measure peer effects by exploiting random assignment to peer groups; in contrast, we study peer effects in naturally occurring peer groups, but exploit random variation in the “price” of a social program for a subset of individuals. Fathers of children born after April 1, 1993 in Norway were eligible for one month of governmental paid paternity leave, while fathers of children born before this cutoff were not. There is a sharp increase in fathers taking paternity leave immediately after the reform, with take up rising from 3% to 35%. While this quasi-random variation changed the cost of paternity leave for some fathers and not others, it did not directly affect the cost for the father’s coworkers or brothers. Therefore, any effect on the brother or the coworker can be attributed to the influence of the peer father in their network. Our key findings on peer effects are four-fold. First, we find strong evidence for substantial peer effects of program participation in both workplace and family networks. Coworkers and brothers are 11 and 15 percentage points, respectively, more likely to take paternity leave if their peer father was induced to take up leave by the reform. Second, the most likely mechanism is information transmission about costs and benefits, including increased knowledge of how an employer will react. Third, there is essential heterogeneity in the size of the peer effect depending on the strength of ties between peers, highlighting the importance of duration, intensity, and frequency of social interactions. Fourth, the estimated peer effect gets amplified over time, with each subsequent birth exhibiting a snowball effect as the original peer father’s influence cascades through a firm. Our findings demonstrate that peer effects can lead to long-run equilibrium participation rates which are substantially higher than would otherwise be expected.

Keywords: Social Interactions, Peer Effects, Program Participation

JEL codes: D62, J13, I38

*Department of Economics, UC San Diego, e-mail: gdahl@ucsd.edu

[†]Department of Economics, University of Bergen, e-mail: katrine.loken@econ.uib.no

[‡]Department of Economics, University College London, e-mail: magne.mogstad@gmail.com

1 Introduction

Economists and policymakers are keenly interested in understanding the effects of social interactions on individual behavior. One question of particular interest is how peer groups influence the take-up of government social programs. Peer groups could serve as important information transmission networks or be influential in changing social norms, particularly in settings where information is scarce and perceptions are in their formative stage. Social interactions could reinforce or offset the direct effects on take-up due to a program's parameters, leading to a long-run equilibrium outcome which is substantially lower or higher than otherwise expected.

Estimating the causal effect of social interactions has proven difficult given the well-known problems of reflection, correlated unobservables, and endogenous group membership. On top of these identification issues, it is often challenging to define the appropriate peer group and access data which links members of a peer group together. Early and ongoing research attempts to control for as many group characteristics as possible or use instrumental variables.¹ Recognizing that estimates could still be biased, another set of papers attempts to measure peer effects by exploiting exogenous assignment to peer groups.²

In contrast, we study peer influence in naturally occurring peer groups, but exploit variation in the “price” of a social program for a random subset of individuals in the spirit of Moffitt's (2001) “partial-population” identification approach. This approach takes advantage of the fact that treatment (i.e., price) is randomly assigned and therefore unrelated to any other factors which might influence take-up.³ As we discuss later, with random variation in treatment (and with group membership determined prior to treatment), the triple threats of reflection, correlated unobservables, and endogenous group membership no longer bias the estimates of peer effects.

¹For examples, see Bandiera and Rasul (2006), Bayer, Ross, and Topa (2008), Bertrand, Luttmer, and Mullainathan (2000), Case and Katz (1991), Carrell et al. (2008), Gaviria and Raphael (2001), Glaeser, Sacerdote, and Scheinkman (1996), Hensvik and Nilsson (2010), Markussen and Roed (2012), Maurin and Moschion (2009), Munshi (2003), and Rege, Telle, and Votruba (2009).

²See, for example, Babcock, Bedard, Charness, Hartman, Royer (2011), Bandiera, Barankay, and Rasul (2009, 2010), Carrell, Fullerton, and West (2009), Carrell and Hoekstra (2010), Carrell, Hoekstra, and West (2011), Cullen, Jacob, and Levitt (2006), Duncan et al (2005), Hanushek et al (2003), Hoxby (2000), Imberman, Kugler, and Sacerdote (forthcoming), Jacob (2004), Katz, Kling, and Liebman (2001), Kling, Liebman, and Katz (2007), Kling, Ludwig, and Katz (2005), Kremer and Levy (2008), Lefgren (2004), Ludwig, Duncan, and Hirschfield (2001), Ludwig et al (2008, 2011), Mas and Moretti (2009), Sacerdote (2001), Stinebrickner and Stinebrickner (2006), and Zimmerman (2003).

³A small but growing literature uses the partial population approach and quasi-experimental variation in treatment to estimate peer effects in naturally occurring, self-chosen social networks. See Angelucci et al (2010), Baird et al (2012), Bobinis and Finan (2009), Bursztyn et al (2012), Duflo and Saez (2003), Kremer and Miguel (2007), Kuhn et al (forthcoming), and Lalive and Cattaneo (2009). None of these studies look at peer effects in participation in social programs.

We estimate peer effects in the context of paid paternity leave in Norway using a regression discontinuity (RD) design.⁴ We study whether social interactions matter for paternity leave take-up along two dimensions: workplace networks (coworkers) and family networks (brothers). Fathers of children born after April 1, 1993 in Norway were eligible for one month of governmental paid paternity leave, while fathers of children born before this cutoff were not. Before the introduction of this paternity leave program, parents had a shared leave quota which could be split between the mother and father. In practice, however, mothers took the entire amount of leave, with very few fathers taking any leave at all. To encourage more fathers to take leave, the 1993 reform stipulated this extra month of paid leave could only be taken by fathers.

There is a sharp increase in fathers taking paternity leave immediately after the reform, from a pre-reform take up of 3% to a post-reform take up of 35%. This quasi-random variation changed the cost (or price) of paternity leave for some fathers and not others. However, it did not directly affect the cost of taking leave for the father’s coworkers or brothers, since they were all eligible for paid paternity leave when they had children in the post-reform period. Therefore, any effect on the coworker or on the brother can be attributed to the influence of the reform-window father in their network (the peer father), and not a change in the fundamental parameters of the leave program.

Our key findings on peer effects are four-fold. First, we find strong evidence for substantial peer effects of program participation in both workplace and family networks. To study the effect of peers in the workplace in an RD framework, we focus on small to medium sized firms where there is a single birth among all employees in the 12-month window surrounding the reform. This window restriction solves several challenges which otherwise make identification of peer effects difficult in an RD framework. We find that coworkers are 3.7 percentage points more likely to take paternity leave if their colleague was eligible versus not eligible for paternity leave around the reform cutoff. Since the first-stage estimate on take up is 32%, this implies a sizeable peer effect estimate of roughly 11 percentage points. For the family network, we find that brothers of reform-window fathers who were eligible for leave are 4.8 percentage points more likely to take paid leave after the birth of their first child. This implies a sizeable peer effect estimate of roughly 15 percentage points. The results for both the brother and coworker networks are statistically significant and robust to a variety of alternative RD specifications and control variables.

⁴To our knowledge, RD has not previously been used to estimate peer effects within naturally occurring peer groups. As discussed below, using an RD approach for this purpose involves a unique set of challenges because multiple peers in a network can affect the same individual.

Second, the most likely mechanism for the peer effect is information transmission about costs and benefits of participation, including how employers will react and whether there is a social stigma. The peer effect operates primarily through the peer father's take up, as we find no statistical evidence for other direct effects on peer fathers as a result of the reform, including no impact on fertility, employment, or earnings. Because variation in the cost of paternity leave near the reform cut-off is as good as random, the peer effects estimates are not picking up common time effects such as general changes in societal norms. Since the parental leave system is universal, simple, and already well-known, the mechanism for the peer effect is not information about either the existence of the program or how to sign up for the program. The mechanism is also not leisure complementarities or direct consumption externalities since coworkers and brothers do not take leave at the same time as the original peer father. Interestingly, we find suggestive evidence that the workplace and family networks transmit different types of information about the costs and benefits of participation. This makes sense, as a coworker can reveal important information about how a particular firm will react, while a brother is more likely to pass on information related to the family setting.

Third, we find essential heterogeneity based on the strength of interpersonal ties between peers. We operationally define the strength of ties by the nature of the relationship and the type of interactions. Strong peer effects are found for long-term familial relationships such as brothers and for male coworkers who have frequent interactions in a firm. We find larger peer effects when ties are arguably stronger, including larger effects in smaller firms and for coworkers who start working at the firm around the same time. Looking at weaker ties in extended family and extended workplace networks, we find no evidence for peer effects. In particular, peer fathers do not appear to influence their brother-in-laws or their female coworker's husbands. We next look at neighborhood peers, defining neighbors precisely (e.g., the two closest households). In our setting of paternity leave, neighbors defined by geography exert no influence on each other. Our findings highlight the importance of duration, intensity, and frequency of social interactions in understanding how peer groups influence the take up of social programs.

Fourth, we find the estimated peer effect gets amplified over time within a firm, with each subsequent birth exhibiting a snowball effect in response to the original reform. The peer effect cascades through the firm network as the first peer interacts with a second peer, the second peer interacts with a third peer, and so on. The total peer effect can be decomposed into the direct influence of the peer father and the indirect snowball effects operating through the increase in take up of other coworkers.

The snowball portion is large, accounting for over 50 percent of the total peer effect for the third and higher-order coworkers in a firm who have a child after the original peer father. We further decompose these direct and indirect effects over time. In the early years after the reform, most of the estimated peer effect can be attributed to the direct effect, as there is little opportunity for intervening births to create a snowball effect. However, over time, the direct influence of the original peer father decays so that later in the sample period it is virtually zero. In contrast, the snowball effect gets larger and larger over time as more coworkers have a child within a given firm. Even though the snowball effects also decay, the accumulation of effects from intervening coworkers more than offsets this decay.

Taken together, our results have important implications for the peer effects literature and for the evaluation of social programs. Our study points out that an individual's choice of participation in a social program can be affected by more than one peer group and that these social interaction effects can be sizeable. Both the workplace and family can serve as important information transmission networks in settings where information about the benefits and costs of program participation is scarce and perceptions are in their formative stage. Our findings highlight that peer effects can have long-lasting effects on program participation, even in the presence of decay, since any original peer effect cascades through a network over time. This is especially important when considering the design and implementation of new social programs, since the initial group of participants can play a large and lasting role in the evolution of take up patterns. Social interactions can reinforce the direct effects on take up due to a program's parameters, leading to a long-run equilibrium take-up rate which can be substantially higher than in the absence of peer effects.

The remainder of the paper proceeds as follows. Section 2 discusses the challenges in estimating social interaction effects, the previous literature, and our identification strategy. In Sections 3 and 4, we discuss the 1993 leave reform, our data, the RD design, and validity tests. Section 5 presents our main findings on peer effects in the workplace and family networks. Section 6 explores possible mechanisms, Section 7 examines the importance of strong versus weak ties, and Section 8 estimates how peer effects cascade through the social network. The final section offers some concluding remarks.

2 Identifying Social Interactions

2.1 Threats to Identification

A social interaction or peer effect occurs when the action of one individual affects the actions of other individuals in the same social group. As Manski (1993) and others have pointed out, estimation of these effects is difficult given the problems of simultaneous causality (reflection), correlated unobservables (contextual effects), and endogenous group membership. To illustrate these identification issues, consider a model which is linear in the social interaction effect. For simplicity, we assume there are only 2 individuals in each group, although this could easily be generalized. Letting y_{ig} denote the outcome for individual i in group g , the system of simultaneous equations for peer effects is:⁵

$$y_{1g} = \alpha_1 + \beta_1 y_{2g} + \gamma_1 x_{1g} + \tau_1 x_{2g} + \theta_1 w_g + e_{1g} \quad (1)$$

$$y_{2g} = \alpha_2 + \beta_2 y_{1g} + \gamma_2 x_{2g} + \tau_2 x_{1g} + \theta_2 w_g + e_{2g} \quad (2)$$

where x_{ig} are observable characteristics of individual i in group g , w_g are characteristics which vary only at the group level, and e_{ig} is an error term. This model captures the idea that individual 2's choice is influenced by the choice individual 1 makes, and visa versa. It also allows individual 2's choice to depend on his own characteristics, the characteristics of individual 1, and common group-specific variables.

The equations above are an example of simultaneous causality bias, since individual 1's choice affects individual 2's choice, and there is no exclusion restriction. Manski (1993) points out that the coefficients are not identified and labels this the reflection problem. The problem of correlated unobservables arises when not all relevant group-level (w_g) or individual variables (x_{1g} , x_{2g}) are observed, leading to an omitted variable bias in the estimated peer effect due to what Manski calls contextual effects. Finally, the problem of endogenous group membership arises when individuals chose which group to belong to as a function of the characteristics and choices of the group.

All three of these problems arise when trying to estimate the take-up of social programs. In our setting which looks at paternity leave in workplace and family

⁵Manski's formulation of the problem replaces x_{ig} and y_{ig} with their expected group values; we use notation similar to Moffitt (2001), since this seems more natural in our setting (2 members in a peer group) and is more general in that it allows for e_{2g} to affect y_{1g} (and e_{1g} to affect y_{2g}). A linear model guarantees a unique equilibrium, rather than multiple equilibria, which may partly explain why it is the most widely used model to study social interactions. Allowing for more than two members in a peer group does not change the key insights; in practice, when there are multiple members of a peer group, most researchers assume a linear-in-means model, where an individual's choice depends on the leave-out mean for the other members in the group.

networks, coworkers and brothers are likely to influence each other. There are also a variety of workplace and family characteristics, such as a family-friendly work environment or supportive grandparents, which are likely to be both unobserved and correlated within groups. While endogenous group membership is less of an issue for brothers, it is an obvious problem for coworkers, as fathers who are inclined to take paternity leave might naturally be attracted to seek employment with coworkers who feel the same way.

2.2 Previous Research

The existing literature has tried several approaches to overcome the challenges inherent in estimating peer group effects. A large set of papers document correlations in behavior and choices within peer groups for a variety of outcomes. The most common research design controls for a large number of group and individual level characteristics in an attempt to minimize the bias caused by simultaneous causality, correlated unobservables, and endogenous group membership. Several studies have also taken care to define networks precisely, narrowing in on the most likely peer group while controlling for more aggregate group effects.

A leading example of this approach for program participation is the study of welfare take up by Bertrand, Luttmer, and Mullainathan (2000). They use language spoken at home and geographical neighborhoods to define peer groups, which allows them to include local area and language group fixed effects. Other examples using this strategy include studies on peer effects in crime (Glaeser, Sacerdote, and Scheinkman, 1996), employment (Bayer, Ross, and Topa, 2008; Munshi, 2003), entrepreneurship (Nanda and Sorensen, 2010), fertility (Hensvik and Nilsson, 2010), disability pension participation (Rege, Telle, Votruba, 2009; Markussen and Roed, 2012), cheating (Carrell et al, 2008), and risky teenage behavior such as drug, alcohol, and cigarette use (Case and Katz, 1991; Gaviria and Raphael, 2001). These studies look at a variety of peer groups, including networks defined by families, neighborhoods, classmates, and coworkers.

The hope in these non-experimental studies is that any remaining bias after carefully controlling for covariates is small. However, some researchers have pointed out inherent difficulties. In an early paper, Evans, Oates, and Schwab (1992) illustrate how endogenous group membership is likely to bias estimates in observational studies. As Ross (2009) argues, neighborhoods are stratified by racial, ethnic, and economic groups, which makes social interaction studies based on geography intrinsically difficult. Currie and Aizer (2004) provide an example where observational estimates, even with an extensive set of controls, pick up neighborhood effects rather than

peer effects. To partly address such concerns, some researchers provide tests of the identifying assumptions made in observational studies (e.g., Bayer, Ross, and Topa, 2008; Hensvik and Nilsson, 2010) or use instrumental variables (e.g., Case and Katz, 1991; Maurin and Moschion, 2009; Monstand, Propper, Salvanes, 2011; Rege, Telle, and Votruba 2009). While these papers represent important contributions to our understanding of peer effects, it is difficult to be certain the assumptions hold or the instruments are valid.

To avoid the problems associated with observational studies, some researchers have taken advantage of random assignment to peer groups. A classic example using random assignment is provided by Sacerdote (2001), who examines peer effects among college roommates. Since freshman roommates are randomly assigned, peer effects can be estimated without worrying about endogenous selection into peer groups or correlated unobservables.⁶ Carrell and coauthors take advantage of exogenous assignment in the military, where individuals are randomly matched with a larger peer group with whom they spend a majority of their time (Carrell, Fullerton, and West, 2009; Carrell and Hoekstra, 2010; Carrell, Hoekstra, and West, 2011).⁷ Random assignment to groups has also been used to study neighborhood effects in a series of influential papers by Katz, Kling, and Liebman (2001), Kling, Liebman, and Katz (2007), Kling, Ludwig, and Katz (2005), Ludwig, Duncan, and Hirschfield (2001), Ludwig et al (2008, 2011). Using lottery assignment of housing vouchers in the Moving to Opportunity program, they examine how neighborhoods affect a variety of short and long run outcomes.⁸

Studies looking at random assignment to peer groups are both convincing and important. They answer the question of what happens when individuals are placed into social networks or environments which are different from what they are used to. But they cannot answer questions about social interactions in naturally occurring, self-chosen peer groups. This distinction is particularly important for designing and evaluating social programs, where a key question is whether endogenously-formed social networks transmit information or otherwise influence participation. As Carrell, Sacerdote, and West (2012) show, endogeneous sorting into natural peer groups is a powerful force. Even when individuals are randomly assigned to modestly-sized peer

⁶Subsequent papers have used random assignment of roommates to study a variety of outcomes for different populations. See Duncan et al (2005); Kremer and Levy (2008); Stinebrickner and Stinebrickner (2006); and Zimmerman (2003).

⁷Related papers use arguably exogenous variation in the assignment of students to classrooms to study academic achievement (Cullen, Jacob, and Levitt, 2006; Hanushek, Kain, Markman, and Rivkin, 2003; Hoxby, 2000; Imberman, Kugler, and Sacerdote, forthcoming; Lefgren, 2004).

⁸A related paper is Jacob (2004), which uses a natural experiment to look at neighborhood effects for displaced residents of housing projects that were demolished in Chicago.

groups, they self-select into more homogeneous sub-groups, subverting the intended peer group assignment.

2.3 Using Experimental Variation within Naturally Occurring Peer Groups

In contrast to the previous literature discussed above, we study naturally occurring peer groups, but exploit variation in the “price” (or cost) of a social program for a random subset of individuals within groups. Instead of randomly assigning individuals to groups and seeing how participation is affected, we randomly vary the net benefit of participation for some individuals in a group and see how other members in the group change their behavior. Moffitt (2001) calls this the “partial-population” approach.

To fix ideas, consider an experiment where (i) the price, p_{1g} , of program participation for individuals with the label 1 is varied randomly across groups and (ii) there is no change for any individuals with the label 2. Equations (1) and (2) become:

$$y_{1g} = \alpha_1 + \beta_1 y_{2g} + \gamma_1 x_{1g} + \tau_1 x_{2g} + \theta_1 w_g + \lambda p_{1g} + e_{1g} \quad (3)$$

$$y_{2g} = \alpha_2 + \beta y_{1g} + \gamma_2 x_{2g} + \tau_2 x_{1g} + \theta_2 w_g + e_{2g} \quad (4)$$

Since p_{1g} is assigned randomly to individuals with the label 1 in group g , it will be uncorrelated with x_{1g} , x_{2g} , w_g , e_{1g} , and e_{2g} . This immediately implies that λ can be identified from a regression of y_{1g} on p_{1g} . More importantly, it also means that a consistent estimate of the peer effect β can be obtained by regressing y_{2g} on p_{1g} and scaling by $\hat{\lambda}$.⁹

The presence of an excluded variable which appears in individual 1’s outcome equation but not individual 2’s solves the reflection problem of simultaneity. Moreover, since p_{1g} is orthogonal to all observed and unobserved covariates, correlated unobservables can no longer bias the estimates. And finally, as long as peer groups are measured before the price shock p_{1g} , endogenous group membership does not create a bias either; any changes in group membership which happen after the price shock are either a causal result of changes in p_{1g} or orthogonal to changes in p_{1g} .

A handful of researchers have used a partial population approach to study social interactions. For example, Kremer and Miguel (2007) study deworming technology adoption among self-identified peers in Kenya; Lalive and Cattaneo (2009), Angelucci et al (2010), and Bobinis and Finan (2009) all look at school attendance within Mexican villages as a result of the Progressa program; and Duflo and Saez (2003)

⁹Since p_{1g} is orthogonal to all other covariates, for consistency it does not matter whether other covariates are included in either regression. With or without covariates, the experimental estimate of β can be interpreted as an IV estimate.

look at attendance at an investment information fair and subsequent retirement investments among colleagues.¹⁰ These papers answer important questions about peer effects, and find the indirect benefits due to spillovers within social networks to be almost as large as the direct effects. In each of these examples, a large portion of the estimated effect is due to consumption or outcome complementarities: when an individual is treated for worms, the benefit of the technology to their peers declines; when a child goes to school, the benefit to their friends of attending school also goes up; and when an employee attends an information session in response to a cash incentive, untreated colleagues in the same peer group now have someone to attend the information seminar with.¹¹

Our study complements this strand of the peer effects literature in several ways. None of the above studies look at peer effects in participation in social programs. Peer effects in program participation are highly policy relevant since they can reinforce or offset the direct effects on take-up due to a program's parameters, leading to a long-run equilibrium outcome which is substantially lower or higher than otherwise expected. Our setting is also fundamentally different, in that there is little role for consumption or outcome complementarities. The reason is that fathers take leave for a limited amount of time and the period when reform-window fathers can take leave is far removed from when their brothers and colleagues are eligible for leave. Since the births are temporally distant, there is no externality due to peers taking leave at the same time. Rather than leisure complementarities, the peer effects we estimate are most likely capturing the transmission of information about the costs and benefits to taking leave, including whether there is a social stigma and how employers will react. As a recent working paper by Bursztyn et al (2012) shows, information channels can play a large role among peers, even in the absence of consumption or outcome complementarities.

Another difference is that we study heterogeneity in the size of the peer effect within naturally occurring social networks depending on the strength of ties between peers. Due to the richness of Norwegian registry data, we are able to demonstrate the importance of duration, intensity, and frequency of social interactions. Lastly, we show how the total peer effect estimate can be decomposed into a direct peer influence and an indirect snowball effect that accumulates within a network over

¹⁰A related paper using the partial population design is Kuhn et al (forthcoming), which looks at neighbors of lottery prize winners in Holland.

¹¹Baird et al (2012) discuss how to design experiments to measure these spillover effects, and apply their framework to a cash transfer program in Malawi. When individual treatment interventions benefit the entire peer community due to spillovers, Angelucci and Di Georgi (2009) argue that experimental designs should use group-level randomization, rather than selecting treatment and control subjects randomly within groups.

time. The decomposition reveals that the snowball effect plays a key role for the evolution of program participation because peer effects cascade through a network as the first peer interacts with a second peer, the second peer interacts with a third peer, and so on.

2.4 Estimating Peer Effects in an RD design

In our setting of paternity leave, the price of leave changes discontinuously based on when a child is born: fathers of children born after April 1, 1993 in Norway were eligible for one month of governmental paid paternity leave, while fathers of children born before this cutoff were not. Using an RD design, we get quasi-random variation in the cost of taking leave for fathers (individual 1 in group g) whose children are born in a window surrounding the reform. We can estimate λ in equation (3) as the jump in take-up at the reform date cutoff in a first stage RD regression. We can then examine whether this quasi-random variation in cost for father 1 changes the leave taking behavior of father 2. This reduced form RD estimate can be scaled by $\hat{\lambda}$ to get an estimate of $\hat{\beta}$. The details of the reform and the RD procedure are outlined in the next section.

To our knowledge, RD has not previously been used to estimate peer effects within naturally occurring peer groups.¹² Using an RD approach for this purpose involves a unique set of challenges because of what might be called the “many to one” feature which is inherent in peer groups. By many to one, we mean that multiple peers in a network can affect the same individual. For example, in our firm setting, a coworker can potentially be affected by multiple peer fathers.

In an RD design, a window surrounding the cutoff (i.e., the reform date) needs to be specified. Several issues arise when multiple peer fathers appear in the chosen RD window. First is the issue of how to define the running variable when there is more than one peer father in the window. This is particularly problematic when there are some peer fathers before the cutoff and some peer fathers after the cutoff. A second issue relates to functional form. Is a coworker affected by (i) the average number of peer fathers with children born before versus after the cutoff, (ii) the number of peer fathers with children born after the cutoff, or (iii) simply whether any peer father had a child after the cutoff? For instance, if the average number is used, the implicit assumption is that the number of peers doesn’t matter: if there is 1 peer father after the cutoff and 2 peer fathers before, this has the same effect

¹²There are studies using an RD design to exploit quasi-random assignment to peer groups. One example is Ding and Lehrer (2007), which uses regression discontinuities created by the entrance exam cutoff rules to study the impact on child development of assignment to secondary schools with different rankings.

as 3 peer fathers after and 6 peer fathers before. Even if the running variable can be defined and a functional form decided on, a final issue is that for large networks, an RD approach will have little power. The reason is that as the number of peer fathers appearing in the reform window increases, the variation in peer exposure to the reform decreases, since roughly an equal number of peer fathers will give birth before versus after the cutoff. As a consequence, precise RD estimates of peer effects within naturally occurring peer groups require many small networks.

We address these issues by restricting the sample so there is a single birth in the reform window. For the firm network, we restrict the sample of peer fathers to firms which have only one birth to male employees in the one-year interval straddling the reform (6 months on each side of the cutoff). This restriction allows us to cleanly identify a single peer father and use a straightforward RD design. With this approach, it is easy to define the running variable, no additional functional form assumptions are needed, and there is ample variation in peer exposure to the reform.

One implication of our approach is that the estimation sample will be comprised of small- and medium-sized Norwegian firms. The median firm size for workers in our restricted sample is 27 employees, while the median firm size for all workers in Norway is 58 employees.¹³ These small and medium firms are suitable for a study of peer effects, because it is likely that employees in these types of firms interact with each other directly. For the family network, we also restrict the sample of peer fathers to families which have only one brother with a child being born in a window straddling the reform. This restriction is generally not binding, as few families have multiple brothers giving birth in the reform window. Further details on our setting, data, and empirical approach are found in the next section.

3 Background, Data and Empirical Strategy

3.1 Paternity Leave

Governmental paid parental leave has a long history in Norway. In 1977, parents were granted 18 weeks of paid leave. During the 1980s and 1990s, the leave period was gradually expanded, and by 2011 there was a maximum of 47 weeks of paid leave. The parental leave mandates offer employment protection and income replacement. The parental leave policy is part of the broader Social Security System, and is financed through employer- and employee-paid taxes. Apart from a few weeks reserved for

¹³We do not have many large firms in our sample, since these firms are likely to have more than one birth in a one year window of the reform. Likewise, we have few of the smallest firms, since these firms are less likely to have any births in a one-year window of the reform.

the mother, parents could share the parental leave between them as they desired before 1993. Until recently, however, fathers were taking little, if any, leave.

To induce fathers to take parental leave, the labor party government introduced a paternal-leave taking quota in their suggestion for the national budget of 1993. The reform was passed in parliament in December 1992 and implemented on April 1, 1993. The key feature of the paternal quota was that four out of 42 weeks of paid parental leave were reserved exclusively for the father.¹⁴ With few exceptions, the family would lose these four weeks of paid parental leave if not taken by the father. Apart from exclusive quotas of four weeks for fathers and the pre-existing nine weeks for mothers, parents could share the parental leave between them as they desired.

While paid maternity leave was only contingent on the mother working at least 6 of the last 10 months before birth, paid paternity leave was contingent on both parents working at least 6 of the last 10 months.¹⁵ Income payments were based on the earnings of the person on leave, but a father's payment was reduced proportionally if the mother did not work full-time prior to birth. In families with full-time working mothers prior to childbirth, the parental leave scheme offers 100 percent income compensation, subject to a capped amount, for both men and women.¹⁶ The income cap is non-binding for most parents.¹⁷ During both maternity and paternity leave, the firm is not allowed to dismiss the worker, and the parent has the right to return to a job that is comparable to the job he or she held before going on leave.

The parental leave system is universal, simple, and well-known (including details about eligibility, benefit amounts, and the application process). To apply for parental leave benefits, parents must inform their employers and submit a *joint* application to a Social Security Administration field office. For each spouse, the family must

¹⁴At the same time as the four-week paternity quota was implemented, the leave amount that could be shared between parents was extended by three weeks. This means that we cannot tell for sure whether the estimated peer effects should be interpreted as reflecting the introduction of the paternity quota or the extension of shared parental leave. We expect, however, that the paternity quota is the driving force behind the jump in the fraction of fathers taking leave after April 1, 1993. The reason is that none of the previous extensions of shared parental leave increased the take-up of leave among fathers.

¹⁵Cohabitation is common in Norway, and the rules allow cohabiting couples to fully participate in the parental leave program.

¹⁶There are two exceptions. First, the replacement rate for self-employed individuals is 65% of income. Second, parents could choose 80% income replacement and receive an additional 6 weeks leave.

¹⁷Benefits are capped at an amount which is six times the "substantial gainful activity" threshold; this threshold is set annually by the government and adjusted in accordance with the average wage growth in the economy. In 2010, the substantial gainful activity level was NOK 72,900 per year (approximately \$12,500) which meant benefits were capped at NOK 437,400 (approximately \$75,000). Thirty-four percent of fathers and only 7 percent of mothers earned more than the benefits cap. In such situations, most employers (private and public) top up the benefits so that income is 100% compensated.

specify days of leave and when the leave period will start and end.¹⁸ The application must be submitted at least six weeks before the pregnancy due date. If the applicant meets the eligibility criteria, a parental leave award is made. Because almost all eligible women take leave and the family must specify maternity and paternity leave on the same form, the introduction of the paternal-leave taking quota had few, if any, practical implications for the application process. The key change was that more families filled in non-zero days of paternity leave in the application form, instead of leaving it blank.

The introduction of the paternity quota led to a sharp increase in take-up rates of parental leave by fathers. Figure 1 shows the fraction of fathers taking paternity leave by the birth year of their child. There is a stark increase in the share of fathers taking paternity leave immediately after the reform: While only 3 percent of fathers took leave prior to the introduction of the paternal quota, the take-up rate jumped to approximately 35 percent in April 1993 after the reform was implemented. The take-up rate continued to rise over the next decade, climbing to 70 percent of eligible fathers by 2006.

3.2 Data and Sample Restrictions

Our analysis employs several data sources that we can link through unique identifiers for each individual. The data on parental leave comes from social security registers that contain complete records for all individuals for the period 1992-2006.¹⁹ The data set contains information on the date an individual was awarded leave, the number of days of paid leave, and the level of benefits received. We link this data with administrative registers provided by Statistics Norway, using a rich longitudinal database that covers every resident from 1967 to 2006. For each year, it contains individual demographic information (including gender, date of birth, and marital status), socio-economic data (including years of education and earnings), and exact geographical identifiers (including street address and zip code). The data contains unique identifiers that allow us match spouses and parents to their children. Lastly, we merge these data sets with linked employer-employee data that contains complete records of all firms and workers for the period 1992 to 2006. A number of firm specific variables are available, such as firm size and industry. The coverage and reliability of Norwegian registry data are considered to be exceptional, as illustrated by the fact that they received the highest rating in a data quality assessment conducted by

¹⁸Most fathers taking leave have only one spell of paternity leave. Their leave period typically comes after the maternity leave period, when the child is (at least) nine months old.

¹⁹Note that since the data on parental leave is not available prior to 1992, a study similar to ours could not be done for earlier parental leave reforms.

Atkinson et al. (1995).

We focus on whether social interactions matter for paternity leave take-up along two dimensions: workplace networks (coworkers) and family networks (brothers). For both peer groups, we restrict the sample to fathers predicted to be eligible in order to gain precision in the RD estimation. Actual eligibility is based on (i) both the father and the mother working at least 6 out of 10 months immediately preceding the birth, and (ii) both the father's and mother's earnings in the prior 10 months exceeding the "substantial gainful activity" threshold. The earnings threshold is set annually by the government and adjusted in accordance with the average wage growth in the economy. In 2010, the substantial gainful activity level was NOK 72,900 per year (approximately \$12,500).

Since we do not observe months of work, we predict eligibility based on earnings in the year prior to childbirth; we count a father as eligible if both the father's and mother's annual earnings exceed the "substantial gainful activity" level. There is a tradeoff between using too strict of an earnings requirement and excluding parents from our sample who were, in fact, eligible, and using a less strict earnings requirement and including parents who actually were ineligible. While including ineligible fathers may increase the residual variation and thus the standard errors in the RD estimation, excluding eligible fathers may affect the external validity of our results.²⁰ By using a fairly weak earnings requirement in the prediction of eligibility, we assign more weight to the generalizability of our results. This conservative approach yields an average take-up for predicted eligible fathers of 60% over the entire post-reform period, while it is only 4 % for predicted non-eligible fathers.

For each peer group, we further refine the sample to be appropriate for the relevant social network. For the family network, we include fathers with a child born of any parity within one year of the reform, who have brothers whose first child is born after the peer father's child and after the reform. For the employment network, we restrict the sample to firms which have only one birth of any parity to male employees in the one-year interval straddling the reform and coworkers whose first child is born after the peer father's child and the reform. As discussed in section 2.4, this restriction allows us to cleanly identify a single peer father and use a straightforward RD design. Note that for brothers we use a window of one year on each side of the reform as our baseline sample. For firms, since we have more observations, we use a smaller window of six months on each side of the reform.

²⁰As long as eligibility cannot be manipulated, the internal validity of the RD estimates are unaffected by the exclusion of ineligibles. Because of the timing of the reform announcement, there is little chance for eligibility manipulation. Indeed, as we show in the next section, we find no statistical evidence for manipulation.

In Appendix Table A1, we document how these sample restrictions which enable an RD analysis affect the average characteristics of fathers in each of our networks. Our coworker sample contains approximately 20 percent of the entire population of eligible fathers with births in the relevant one-year window. Because of our sample restrictions, the reform window fathers in our employment network sample are on average less educated and slightly less likely to be married, but have otherwise similar characteristics for age, child gender, and number of children. Our brother sample contains approximately 13 percent of eligible fathers with births in the corresponding two-year window. The reform window fathers in our brother network sample are younger on average and as a consequence, also less likely to be married, but are similar with respect to college education, child gender, and number of children.

Our primary variable of interest is the take-up of parental leave by fathers, a dummy variable which equals one if a father takes any amount paid parental leave. We also consider the number of days of paid leave taken by fathers. In addition, we look at several labor market outcomes, including annual gross earnings (which includes wages and income from self-employment), employment status (defined by whether an individual’s earnings exceeds the “substantial gainful activity” level discussed above), and family income (incorporating annual earnings and cash transfers less taxes). We also consider the following family outcomes: marriage, divorce, and fertility. The last outcome that we consider is the exam scores of children at the end of compulsory lower secondary school; these scores are important for admission to upper secondary school. Grades take integer values from one to six and are standardized to be mean zero and standard deviation one for ease of interpretation.

We will estimate social interaction effects with and without a set of pre-determined covariates. These covariates do not significantly influence the results, as expected with an RD design. We document that the covariates are balanced on the two sides of reform cutoff in Table A2, which will be discussed later.

3.3 Empirical strategy

We use a fuzzy RD design to estimate the peer effects of parental leave take up. The discontinuity we exploit arises from the introduction of the paternity quota: fathers of children born after April 1, 1993 were eligible for paid paternity leave, while fathers of children born before this cutoff were not. There is a sharp increase in fathers taking paternity leave immediately after the reform, from a pre-reform take up of 3% to a post-reform take up of 35%. This quasi-random variation changed the cost of paternity leave for some fathers and not others. However, it did not directly affect the cost of taking leave for the father’s coworkers or brother, since they were

all eligible for paid paternity leave when they had children in the post-reform period. Therefore, any effect on the coworkers or brother can be attributed to the influence of the peer father in their network, and not a change in the fundamental parameters of the leave program.

For a given network, the fuzzy RD estimand of the peer effect can be written as

$$\hat{\beta} = \frac{\lim_{\epsilon \downarrow 0} E(y_{2g}|t = c + \epsilon) - \lim_{\epsilon \uparrow 0} E(y_{2g}|t = c + \epsilon)}{\lim_{\epsilon \downarrow 0} E(y_{1g}|t = c + \epsilon) - \lim_{\epsilon \uparrow 0} E(y_{1g}|t = c + \epsilon)} \quad (5)$$

where t indicates the date of birth of the child of individual 1 (peer father) in peer group g (firm or family), the cutoff date c equals April 1, 1993, y_{i1} is an indicator variable equal to one if individual 1 takes up the paternity leave, and y_{2g} is an indicator variable equal to one if individual 2 (peer father's coworker or brother) takes up paternity leave; in the empirical analysis, we will also consider a number of other outcome variables.

In other words, the peer effect is identified by dividing the jump in the take-up rate of paternity leave of the peer father's coworker or brother at c by the jump in the fraction of peer fathers that take up paternity leave at c . The identifying assumption of our fuzzy RD design is that individuals are unable to precisely control the assignment variable, date of birth, near the cutoff, in which case the variation in treatment at c is random. The concern for strategic timing of births is addressed empirically below.

The RD design can be implemented by the following two-equation system:

$$y_{2g} = \alpha_2 + \beta y_{1g} + 1[t \geq c]f_l(t - c) + 1[t < c]f_r(c - t) + e_{2g} \quad (6)$$

$$y_{1g} = \alpha_1 + 1[t \geq c](g_l(t - c) + \lambda) + 1[t < c]g_r(c - t) + e_{1g} \quad (7)$$

where f_l, f_r, g_l , and g_r are unknown functional forms. We will estimate the system of equations using either polynomial or local linear regressions.

We can estimate λ as the jump in take-up at the reform date cutoff in a first stage RD regression, given by equation (7). By estimating the following reduced form model, we can examine whether this quasi-random variation in cost of paternal leave for the peer father (assigned the label 1) changes the leave taking behavior of the peer father's coworker or brother (assigned the label 2):

$$y_{2g} = \gamma_2 + 1[t \geq c](f_l(t - c) + \pi) + 1[t < c]f_r(c - t) + u_{2g} \quad (8)$$

where π can be interpreted as an "intention-to-treat" (ITT) effect of the paternity quota on the leave taking behavior of the peer father's coworker or brother. The 2SLS

estimate of β gives the peer effect. In the terminology of Angrist and Imbens (1994), the estimate of β should be interpreted as the local average treatment effect (LATE) of the reform-induced increase in the peer father’s leave on the leave taking behavior of his coworker or brother. Since the two-equation system is exactly identified, the 2SLS estimate is numerically equivalent to the ratio of the reduced form coefficient π and the first stage coefficient λ , provided that the same bandwidth is used in equations (7) and (8) in the local linear case, and the same order of polynomial is used for f and g in the polynomial regression case.

4 Potential Manipulation

4.1 Strategic Timing of Births

The validity of our RD design requires that individuals cannot manipulate the assignment variable, which is the birthdate of the peer father’s child. If date of birth cannot be timed in response to the paternity leave reform, the aggregate distribution of the assignment variable should be continuous around the cutoff date; an increase in the density would indicate strategic timing of births.

There is little opportunity to strategically time conception, as the implementation date for the reform was announced less than nine months in advance. The national budget which proposed the paternal quota was publicly introduced on October 7, 1992 and passed by parliament in December of the same year. Therefore, mothers giving birth close to April 1, 1993 were already pregnant before the announcement of the reform. Searches in newspaper archives indicate the date of implementation was not discussed publicly before the national budget was passed. Furthermore, the month of implementation varied for previous parental leave reforms.²¹

While strategic timing of conception is unlikely, it is still possible that mothers with due dates close to the cutoff date could postpone induced births and planned cesarean sections. In contrast to current birth practices in the U.S., the vast majority of births in Norway around the time of the reform were spontaneous vaginal deliveries. In 1993, the fraction of children born by cesarean section was 12 percent, and of these deliveries, 59 percent were emergency operations. On average, 12 percent of vaginal deliveries in 1993 were induced, while 88 percent were spontaneous.²² It is therefore possible that parents strategically delayed birth to qualify for the 4 extra

²¹The implementation dates of the previous parental leave reforms in Norway were July 1 (in 1977, 1988, and 1991), May 1 (in 1987, 1990) and April 1 (in 1989 and 1992).

²²See Folkehelseinstituttet, <http://mfr-nesstar.uib.no/mfr/>. In comparison, in 2009 in the U.S. the c-section rate was 33% and 23% of births were induced (National Vital Statistics Reports, Births: Final Data for 2009, Vol. 60, No. 1, November 3, 2011).

weeks of paternity leave.

We test for strategic timing in Table 1 by regressing the birthdate of the child on dummies for one week intervals before and after the reform date of April 1, 1993. To increase precision, for this regression we use the entire sample of all births in Norway between 1992 and 2006 to fathers eligible for any type of parental leave, and not just births to fathers who have brothers or coworkers. We also control for day of week, month, year, and day of year in the regression. We find some evidence that a small number of births are delayed. In the week immediately before the reform, there are an estimated 10 fewer births (relative to the average of 840 births per week); in the week immediately after the reform, there are an estimated 11 more births. Both of these differences are significantly different from zero. However, we do not find evidence of delay further away from the reform window; this is as expected, since it is medically difficult to delay childbirth for very long and most inductions in Norway were for late-term pregnancies.

To avoid the possibility that some births in our sample are strategically delayed, our baseline RD results exclude the week immediately before and the week immediately following the reform date of April 1, 1993. As we will show, using a wider donut of 2 weeks (2 weeks on each side of the reform), or no donut at all, does not materially affect our findings. Figure 2 graphically illustrates there is no measurable effect on fertility in either the workplace or family network samples with a one-week donut. While there are seasonal patterns in the number of births (with more births in the spring), there is no jump in fertility around the discontinuity.

4.2 Eligibility

Another threat to our identification strategy is that the announcement of the reform could cause a change in eligibility among peer fathers around the cutoff date. If it did, then restricting the sample to eligible fathers could bias the estimated peer effects.

As explained above, we predict eligibility based on annual earnings in the year prior to childbirth; we count a peer father as eligible if both the father's and mother's annual earnings exceed the substantial gainful activity level. As the vast majority of fathers (91%) already work and earn more than the substantial gainful activity level, there is little opportunity for reform-induced changes in eligibility for peer fathers. It is still possible that mothers' labor supply (or earnings) could respond to the announcement of the reform. Recall, however, that predicted eligibility of fathers who have a child (in the window surrounding the reform) is based on annual earnings in 1992. As the reform was announced in December 1992, it leaves the

mother with only one month in which she can increase her earnings enough to make the father eligible by our definition. Given this short time frame, there is limited scope for mothers to manipulate the predicted eligibility status of the father.

Figure 3 graphically illustrates there is no measurable change in predicted eligibility of peer fathers around the cut-off date. While there is some seasonal variation in earnings and thus in predicted eligibility, there is no jump in the fraction of predicted eligible fathers around the discontinuity. In Appendix Table A2, we test whether predicted eligibility is directly affected by the April 1, 1993 reform. We run the RD regression given by equation (7) with predicted eligibility as the dependent variable, and find no statistical evidence of direct effects on predicted eligibility.

4.3 Covariate Balance

If families time date of birth or change eligibility status in response to the reform, then we would expect to see changes in the distribution of pre-determined characteristics of the parents around the reform date of April 1, 1993. In Table A2, we test whether these covariates are directly affected by the 1993 reform. We run the RD regression given by equation (7) with individual, family, and child characteristics as the dependent variable. As in our baseline RD results, we exclude the week immediately before and the week immediately following the reform date. It is reassuring to find that the RD estimates are close to zero and always insignificant.

5 Results for the Coworker and Brother Networks

5.1 Graphical Results

A virtue of the RD design is that it provides a transparent way of showing how the peer effects are identified. To this end, we begin by providing a graphical depiction of how our main outcomes vary around the cutoff date before turning to a more detailed regression-based analysis.

Figure 4 displays the fraction of peer fathers taking leave in a window surrounding the reform. The top graph plots this first stage for the workplace network and the bottom graph for the brother network. In both graphs, the running variable, date of birth, has been normalized so that April 1, 1993 is time zero. For the workplace network (top graph), each observation is the average number of peer fathers taking paternity leave in one-week bins, based on the birthdate of their child. For the brother network (bottom graph), we plot unrestricted means for two-week bins since we have fewer observations.²³ For both networks, there is a sharp jump in the take-up

²³There are 242 brothers (with 233 peer fathers) and 550 coworkers (with 153 peer fathers) on

rate of peer fathers at the cutoff, with program participation rising from around 3% to approximately 35%. These graphs provide strong evidence that the reform had large direct effects on the leave behavior of peer fathers. As we will document later in Section 6.1, the reform had no other direct effects on peer fathers for outcomes we observe, including no direct effects on labor market outcomes, child achievement, marital stability, or fertility.

Figures 5 through 7 capture our main results on peer effects. In each graph, we plot unrestricted averages in one-, two-, or four-week bins and include estimated regression lines using either linear trends applied to each side of the cutoff date or a local linear regression. Whereas the regression lines better illustrate the trends in the data and the size of the jumps at the cutoff dates, the unrestricted means indicate the underlying noise in the data. Each graph sets the scale of the y-axis to ± 3 standard deviations of the respective variable. By standardizing the y-axes in this way, we can easily compare the trends in the data and the sizes of the jumps at the cutoff dates across graphs.

Figure 5 plots coworker’s leave take up as a function of the birthdate of their peer father’s child around the reform window. The jump at the cutoff is the ITT estimate. As a reminder, these coworkers are all eligible for the extra four weeks of exclusive paternity leave since they have their first child after the reform has been implemented. The difference is that some coworkers had peer fathers who were not eligible for 4 extra weeks (those observations to the left of reform, labeled 0 in the graphs) while other coworkers had peer fathers who were eligible (those observations at or to the right of 0). The top graph shows data aggregated to one-week bins while the bottom graph aggregates to two-week bins. Both graphs reveal a sharp jump in leave take-up by coworkers if their peer father had his child immediately after (versus immediately before) the reform date of April 1, 1993. Figure 6 presents similar graphs for the brother network, with two-week bins in the top graph and four-week bins in the bottom graph. Again, there is strong visual evidence of a sizeable reduced form effect of the paternity quota on the leave taking behavior of the peer father’s brothers.²⁴

In Figure 7, we present local linear regression graphs for coworker’s and brother’s

average in a one-week interval. While few fathers have multiple brothers, each peer father has an average of 3.6 coworkers. We use larger bins for the brother network throughout, since our samples contain relatively fewer brothers compared to coworkers. In all the first stage graphs, the scale of the y-axis is set to ± 5 standard deviations of the outcome variable.

²⁴In Figures 5 and 6 there is a negative slope as a function of the running variable, both before and after the cutoff. This negative slope is a function of the sample restriction that coworkers and brothers have their children after their peer father and after the reform cutoff, which affects when coworkers and brothers have children during our sample period. It does not create a problem for consistency, since the effect is continuous through the cutoff.

leave take up. The plotted local linear regression lines are based on daily, individual-level data, using a uniform kernel and a bin width of 120 days for the coworker network and 240 days for the brother network. If anything, the jump at the reform cutoff date is even larger for these local estimates.

In summary, the patterns in Figures 5 through 7 indicate substantial peer effects in both the workplace and family networks. The quasi-random variation in the cost of taking leave for reform-window fathers changes the leave taking behavior of their peer coworkers and brothers. To arrive at estimates of the peer effects for each of these networks, the final step is to divide the jumps in Figures 5 through 7 by the corresponding jumps in graphs like those in Figure 4 (where the corresponding graphs use the same order polynomial or local linear estimator). We provide precise estimates of the size of these jumps, their ratios, and their statistical significance in the following section.

5.2 *Baseline Regression Results*

Having shown the raw patterns of leave taking behavior around the reform cutoff, we now turn to regression-based estimates. Table 2 presents the baseline RD estimates for the peer effects of fathers on their male coworkers and brothers. The specifications use daily data, exclude observations in a one-week window on either side of the discontinuity, include linear trends in birth day on each side of the discontinuity, and employ triangular weights. We also include pre-determined control variables for father's and mother's years of education, father's and mother's age and age squared at birth, parent's county of residence and marital status prior to the birth, and an indicator for the gender of the child.

As a reminder, the workplace sample is restricted to firms which have only one birth of any parity to male employees in the one-year interval straddling the reform. The RD estimates for coworker peer effects use a 6 month window on each side of the reform. Coworkers must have their first child after the peer father's child is born and after the reform. We note that each peer father has on average 3.6 coworkers in our sample. Therefore, we cluster all of the standard errors in panel A at the firm level.

Column 1 of Table 2 estimates the first stages and corresponds to Figure 4. For both the workplace and family network, the estimate is little over 30 percentage points. This is a sizeable direct effect on parental paternity leave, driven by an increase in take-up from roughly 3% to 35%. Before turning to our peer effect estimates, we first formally test whether there is any evidence that fertility was affected by the reform, since this is central to the validity of our RD approach. Indeed, as we visually saw in Figure 2, there is no evidence of a spike in births around

the reform cutoff, with an RD estimate close to zero for both networks.

To estimate the workplace peer effect, in column 2 we first estimate the reduced form (or ITT) effect. The RD estimate of a coworker’s leave take up at the cutoff date for their peer father’s child’s birthdate is 3.7 percentage points, a point estimate which is statistically significant at the 1% level. This estimate corresponds to Figure 5. To convert this into the estimated peer effect, we divide the reduced form coefficient in column 2 by the first stage coefficient in column 1. This yields a second stage estimate (which can be calculated via 2SLS since the system of equations is exactly identified) of 11.5 percentage points. This estimated peer effect is large relative to the average take-up rate of 68% for coworkers of untreated fathers.

In panel B of Table 2, we find strong evidence for peer effects among brothers as well. Brothers of reform-window fathers who were eligible for leave are 4.8 percentage points more likely to take paid leave after the birth of their first child. This implies a peer effect estimate of 15.5 percentage points. This represents a substantial increase in take up due to peer effects given the average take-up rate of 57% for brothers of untreated fathers.²⁵

5.3 Robustness Checks

In this section we probe the stability of our baseline estimates to alternative specifications. We conclude that our estimated peer effects are remarkably robust to the usual specification checks performed in RD studies.

In Table 3, we perform a variety of alternative specifications for both the workplace and family networks. We first exclude all control variables from the regressions, and find virtually no change in the estimates. This is to be expected, since the values of pre-determined covariates should not affect the estimated jump at the cutoff date in a valid RD design. We next explore what happens when we use separate quadratic or cubic trends on each side of the discontinuity, rather than separate linear trends. The estimated reduced form and second stage coefficients are slightly larger, although the cubic trend estimate is no longer significant for the workplace results. The next set of robustness checks estimate RD regressions without a one-week donut around the reform date and with a two-week donut, respectively. The results remain significant, and if anything, get somewhat larger the bigger the donut. We also try a specification which includes all of the predicted non-eligible fathers, which yields similar results compared to our baseline estimates. Finally, note that we have been clustering our standard errors at the firm level or the family level. An alternative is to cluster at the

²⁵The average take-up rate of 68% for coworkers is larger than the average take-up rate of 57% for brothers due to the fact that take up increases over time, and brothers have their children earlier in our sample period compared to coworkers.

level of the running variable, which is the day of birth. This alternative clustering does little to the standard errors.

Table 4 varies the window sizes for our baseline results. For the workplace network, in panel A we find that windows of 3 months, 4.5 months, and 6 months (our baseline) yield similar results which all remain statistically significant. The estimates using a smaller window are somewhat larger, but also have larger standard errors. A similar set of results holds in panel B for the brother network. As a reminder, since we have fewer brothers compared to coworkers, we use wider windows of 6 months, 9 months, and 12 months (our baseline) in panel B. As with the workplace network, estimates for the brother sample using a smaller window are somewhat larger, with larger standard errors, but remain statistically significant.

An alternative approach to using polynomials on each side of the reform cutoff is to use local linear regression. Some researchers find this estimation method to be more robust to trends away from the cutoff point. In Table 5, we estimate local linear regressions for the workplace and family networks with bandwidths of varying size. Whether we use a bandwidth of 60 days, 90 days, or 120 days for the coworker sample, we find statistically significant peer effects of 13 to 14 percentage points. A similar finding of robustness holds for the brother sample when we use bandwidths of 120 days, 180 days, or 240 days.

As a further check, we run a series of placebo tests. To do this, we first assign a window around a false reform date (i.e., a false cutoff), and then use the RD approach described in section 3.3 to estimate a reduced form peer effect. We run 730 placebo tests for each network (2 years of estimates), where each estimate increases the false reform date by one day.²⁶ To avoid having these placebo estimates be influenced by any jump at the true cutoff, the placebo windows start after the true reform date of April 1, 1993. Figure 8 graphs the distribution of placebo estimates for both the workplace and family network. As the graphs make clear, the true peer effect (from Table 2) is more extreme than all of the placebo estimates for brothers and almost all of the placebo estimates for coworkers. These findings indicate the odds of finding peer effects as large as we do merely due to chance are small.

In summary, we find the estimated coworker and brother peer effects are remarkably robust to control variables, order of polynomial, donut size, method of clustering, window size, and alternative local linear estimators with varying bin widths.

²⁶It should be noted the placebo estimates are not independent of one another, since the windows contain a significant amount of overlap in observations.

6 Mechanisms

As long as individuals cannot manipulate the assignment variable, our RD design provides consistent estimates of the reform-induced increase in peer father’s leave on the leave taking behavior of his coworker or brother. To clarify the nature of this relationship, we take two additional steps. We first examine if there are other direct effects of the reform on the peer father besides increased take up which might serve as mediating relationships. We next assess three different channels through which the increase in peer father’s leave may affect the leave taking behavior of his coworker or brother.

6.1 *Mediating Relationships*

In Table 6, we test for a variety of direct effects of the reform on other outcomes, but find no measureable changes in these potentially mediating outcomes. Except for changing the dependent variable, the RD estimates in the table use the same specification as the first stage estimates appearing in Table 2. There is no evidence of a statistically significant discontinuity in the future employment and earnings of fathers or mothers, or in the relative employment and earnings of mothers versus fathers. There is also no evidence of a direct effect on the grade point average of the child in middle school, completed fertility, or long-term marital status. As documented in the table, these estimates are close to zero and never statistically significant. The only estimate which approaches statistical significance is father’s total earnings, but even so the estimated effect is small, amounting to less than a 2% reduction in earnings.²⁷ The lack of other direct effects suggests that mediating relationships do not play an important role in governing the influence of the reform-induced increase in the peer father’s leave on the leave taking behavior of his coworker or brother.

6.2 *Peer Effect Channels*

Because variation in the cost of paternity leave near the reform cut-off is as good as random, the peer effects estimates are not picking up common time effects such as general changes in societal norms. There are, however, several other channels through which the reform-induced increase in a peer father’s leave may affect the leave taking behavior of his coworker or brother.

²⁷This finding is consistent with Rege and Solli (2010). Using a difference-in-differences approach, they find that the paternity leave reform led to a small but statistically insignificant reduction in future earnings for fathers of newborn children in 1993.

The first possible channel is sharing of information about how to enroll in the program.²⁸ As discussed in Section 3.1, the parental leave system is universal, simple, and well-known (including details about eligibility, benefit amounts, and the application process). To apply for parental leave benefits, the spouses must inform their employers and submit a joint application to the government. Because almost all eligible women take leave and the family must specify maternity and paternity leave on the same form, the introduction of the paternal-leave taking quota had few, if any, practical implications for the application process. For these reasons, we do not think a key mechanism for the estimated peer effects is information about either the existence of the program or how to sign up for the program.

The second possible channel is leisure complementarities or direct consumption externalities. Since the births are temporally distant, coworkers and brothers do not take leave at the same time as the original peer father. As a consequence, there is limited scope for leisure complementarities or direct consumption externalities arising from the reform-induced take-up of paternity leave. Another piece of evidence against this channel is that the peer effect is present even if brothers live in different municipalities, as documented in Table 7. This finding is consistent with a recent working paper by Bursztyn et al. (2012) showing that information transmission can play a large role among peers, even in the absence of outcome complementarities.

The third channel is information about the costs and benefits of participation, including how employers will react and whether there is a social stigma. In our setting of paternity leave, information about costs and benefits is scarce and perceptions are in their formative stage. The reason is that prior to the 1993 reform, almost no fathers were taking paternity leave, with the result that few fathers had direct knowledge about the pros and cons of taking leave. However, the reform generates random variation in the take up of peer fathers and therefore changed the information set of a subgroup of peer brothers and coworkers. This exogenous increase in information reduces uncertainty, which should increase take-up among risk averse individuals with unbiased expectations.

Without data on subjective expectations and individual information sets, it is difficult to assess what type of information transmission is driving the estimated peer effects. However, we expect differing pieces of information to be transmitted in the workplace versus family network. In particular, a coworker can reveal important information about the firm-specific consequences of paternity leave, while a brother is more likely to pass on information related to the family setting (or the labor

²⁸Figlio et al. (2011) show that neighborhood social networks can be important in spreading information about eligibility rules and benefits among immigrants with limited knowledge of social programs.

market more broadly). Interestingly, we find several pieces of suggestive evidence which indicate that workplace and family networks indeed transmit different types of information about the costs and benefits of participation.

The first piece of evidence relates to the idea that the informational value about the firm specific consequences of taking leave is likely to be higher if the peer father is a senior manager in the firm. Since we do not have information about the management hierarchy within the firm, we assume the senior managers are the employees with the first or second highest wage in the firm. Table 7 reveals the estimated peer effect is over two and a half times larger if the peer father is a senior manager in the firm as compared to a regular coworker. Another piece of evidence is that coworkers who remain in the same firm have a larger estimated peer effect compared to coworkers who have changed employment by the time they have their child (see Table 7). This is consistent with coworkers learning firm-specific information from the peer father; if the coworker switches firms, the peer father from their old firm has less of an impact on the coworker's leave decision.

We next compare leave take up by type of firm. Workers in the public sector or in firms in highly unionized industries tend to have more secure jobs and regulated pay scales. Consequently, they do not need to worry as much about an employer reacting badly to paternity leave. For workers with high job security, the benefit of learning about a peer father's leave-taking experience should therefore be less valuable. In Table 7, we test whether this is true. Consistent with the job security hypothesis, the estimated peer effects are twice as large in both private sector and low unionization workplaces. We next break up firms based on the average tenure of workers within a firm. In the approximately 25 percent of firms where average tenure is 10 years or more, the estimated peer effect is close to zero. In contrast, for less established firms with higher worker turnover, the peer effect is large and statistically significant. Taken together, these firm-type results suggest the benefit of workplace-specific information is more valuable in settings where there is more job uncertainty.

The final piece of evidence exploits that the perceived productivity signal to the employer is likely to change discontinuously if the father's leave period exceeds the four-week paternity quota. The reason is that the family loses the four weeks of paid paternity leave if not taken by the father, whereas additional days of paternity leave simply crowd out maternity leave. Taking more than four weeks of paternity leave could serve as a signal to employers that a worker is less committed to the job in a way that taking exactly four weeks of leave does not. In Table 8, we estimate peer effects by the duration of the paternity leave spell of the coworker and the brother.

Each row reports separately estimated reduced form effects (columns 1 and 3) and peer effects (columns 2 and 4), using the same specification as in Table 2; the only difference is that the dependent variable in the reduced form and the second stage is now specified as an indicator variable for whether the coworker (columns 1-2) or the brother (columns 3-4) takes more than the given number of weeks of leave. In both networks, the peer effect generates a non-trivial increase in days of leave. But more importantly, we find that brothers are 9.5 percentage points more likely to take more than four weeks of leave because of the reform-induced increase in paternity leave of the peer father; in contrast, there is no evidence of peer effects leading to a crowding out of maternity leave in the workplace network.

7 Strength of Ties

It is natural that some peer groups might exert a stronger influence than others. In a seminal study, Granovetter (1973) classifies interpersonal ties into three categories: strong, weak, or absent. Formally, the strength of ties is defined by the overlap in network members; operationally, the strength of a tie is usually defined by the nature and duration of the relationship as well as by the frequency and intensity of interactions. In this section, we use this operational definition to explore heterogeneity in peer effects based on the strength of interpersonal ties between peers.

While weak ties play an important role in Granovetter's setting, in our setting it is strong ties that are likely to matter most.²⁹ The reason is that prior to the 1993 reform, almost no fathers were taking paternity leave, with the result that few fathers had direct knowledge about costs and benefits prior to the reform. Peers with strong ties are more likely to interact with each other and trust each other's opinions, increasing the chance that information will actually be transmitted and acted upon.

7.1 *Networks with Weaker Ties*

In Section 5, we documented sizeable and robust peer effects among brothers and male coworkers. These peer groups have strong ties as judged by the nature, duration, intensity, and frequency of social interactions. Brothers have known each other for a long time, share a familial bond, and are likely to keep in touch with each other.

²⁹In Granovetter's setting of job finding, he argues for the importance of weak ties since more novel information flows from peers who are part of different social circles. In contrast, he argues that strong ties are less important, since these peers have information sets about available jobs that overlap considerably with what one already knows.

Similarly, male coworkers are likely to have frequent and time-consuming interactions with each other in small to medium sized firms.³⁰

Is there any evidence for peer effects when the ties are weaker? To answer this question, we first turn to extended family and extended workplace networks. In Table 9, we estimate whether a peer father influences his brother-in-law. This tie is arguably weaker than between brothers both in duration and intensity. We find no evidence of any peer effect in this weaker family network. In the second panel of Table 9, we estimate whether a peer father affects his female coworker’s husband. We find no evidence of a significant effect, which is as expected given that a female coworker’s husband is a relatively weak tie who generally works for a different firm.

To explore the peer strength of neighbors, the final panel in Table 9 defines peer groups by geographical neighborhoods. We have the street address of all fathers in Norway, so we define neighborhoods very precisely: we take the two closest households on each side of a father as neighbors. Neighborhoods are defined as of 1993. Similar to the approach used for the firm sample, we limit the sample to “neighborhoods” where there is one birth in a one year window surrounding the reform, and then look at first births to neighbors who had children after the reform and after the peer father. Interestingly, neighbors defined in this way exert no peer influence on each other for paternity take up. This result holds even if we define neighborhoods more broadly; we find similar results using the four closest households or the entire street. Apparently, in this setting, neighbors are not important peers. This does not mean that friends with strong ties exert no peer effect, but rather that neighbors defined strictly by geography seem to have little influence on program participation in our setting.

In each of these weaker networks, the coefficients are small and not significantly different from zero. By way of comparison, the sample sizes and standard errors for the extended workplace and extended family networks are similar to the brother and male coworker networks. So the finding of no significant effect is not due to overly imprecise estimates. Indeed, for the neighborhood network, the sample size is very large and the standard errors are smaller than those for our baseline estimates appearing in Table 2.

³⁰There is little overlap in the two networks. Four percent of coworkers are brothers and 9% of brothers are coworkers in our two samples. While we do not have enough of these observations to estimate the combined peer effect, omitting these observations does not appreciably change our baseline estimates.

7.2 Workplace Networks with Strong versus Weak Ties

In Table 10, we further explore the differential effect of strong and weak ties by focusing on different workplace settings. Coworkers who start working at a firm at a similar time are more likely to be in similar jobs and work with each other. They are also likely to face a similar set of workplace issues related to tenure, such as promotion possibilities. Indeed, we find that workers with start dates at the firm which are within one year of each other are much stronger peers.

Next we look at smaller versus larger firms, since the frequency and number of interactions should vary with firm size. We find the estimated peer effect is approximately 40 percent larger in small firms (less than 27 employees, which is the median firm size in our sample) compared to relatively larger firms. In the final set of comparisons in Table 10, we separate firms by whether they are located in a rural versus an urban area. Interestingly, the estimated peer effect is over twice as big for firms located in rural areas. This partly captures the fact that rural firms are smaller on average. But rural areas also have a much smaller population, so the larger effect may also be driven by the fact that in rural areas coworkers are more likely to be friends outside of the workplace setting.

8 Snowball Effects


Peer effects can play an important role in the evolution of program participation, because peer effects cascade through a network as the first peer interacts with a second peer, the second peer interacts with a third peer, and so on. In our setting, the peer effect could amplify over time within a firm, with each subsequent birth exhibiting a snowball effect in response to the original reform.³¹ This causal chain is initiated by the direct effect of the reform, inducing peer fathers (coworker 1) of children born after April 1, 1993 to take up paternity leave. The second link in the chain is the first subsequent coworker to have a child (coworker 2); his leave behavior is influenced directly by the (reform-induced increase in) leave taking of the peer father. The third link is coworker 3 who has a child after coworkers 1 and 2: the direct influence of the peer father is now amplified by a snowball effect due to the (peer-father-induced increase in) take up of coworker 2. The causal chain continues in this fashion, such that the direct influence of the peer father on coworker i is amplified by a snowball effect operating through the $i - 2$ previous coworkers who had a child.

³¹In principle, the same could happen in the family network. In practice, however, there are too few peer fathers with more than one brother in our sample.

In Section 5 we estimated the total peer effect, which included both the direct influence of the peer father and any indirect snowball effects operating through the increase in take up of other coworkers. The goal of this subsection is to decompose the total peer effect into the direct effect and the snowball effect, and graph the relative importance of these effects over time.

8.1 Identifying Snowball Effects

The following diagram illustrates both the direct effects of a peer father's influence and the indirect, or snowball, effects for the case with four coworkers:



The reform directly influences the peer father, as captured by the horizontal arrow. The peer father directly affects the next coworker after him who has a child, and similarly directly influences coworkers 3 and 4. These direct effects are captured by the bottom arrows in the diagram. But the peer effects of the reform do not stop there. The peer father's effect on other colleagues continues, since coworker 2, who was influenced directly by the peer father, now affects both coworkers 3 and 4. Moreover, coworker 3, who was influenced by both the peer father and by coworker 2 because of the reform, also affects coworker 4. These snowball effects are captured by any path that travels through the top arrows in the diagram.

To make the idea of a snowball effect more precise, we need some notation. Continuing with the case of three coworkers, the causal chain is described by the following system of equations:

$$\begin{aligned}
 y_{1g} &= \alpha + \lambda p_{1g} \\
 y_{2g} &= \alpha_1 + \beta_1 y_{1g} \\
 y_{3g} &= \alpha_2 + \beta_2 y_{2g} + \beta_1 y_{1g} \\
 y_{4g} &= \alpha_3 + \beta_3 y_{3g} + \beta_2 y_{2g} + \beta_1 y_{1g}
 \end{aligned}$$

where the price, p_{1g} , of program participation for the peer father (with subscript label 1) varies randomly across firms (denoted by g), and coworkers are sorted by birth order so that coworker j is the j th father in the firm that has a birth.

Random variation in p_{1g} can be used to identify a set of reduced form coefficients:

$$\frac{dy_2}{dp_{1g}} = \frac{dy_2}{dy_1} \frac{dy_1}{dp_{1g}} = \beta_1 \lambda = \pi_2$$

$$\frac{dy_3}{dp_{1g}} = \frac{dy_3}{dy_2} \frac{dy_2}{dy_1} \frac{dy_1}{dp_{1g}} + \frac{dy_3}{dy_1} \frac{dy_1}{dp_{1g}} = (\beta_2 \beta_1 + \beta_1) \lambda = \pi_3$$

$$\begin{aligned} \frac{dy_4}{dp_{1g}} &= \frac{dy_4}{dy_3} \frac{dy_3}{dy_2} \frac{dy_2}{dy_1} \frac{dy_1}{dp_{1g}} + \frac{dy_4}{dy_3} \frac{dy_3}{dy_1} \frac{dy_1}{dp_{1g}} + \frac{dy_4}{dy_2} \frac{dy_2}{dy_1} \frac{dy_1}{dp_{1g}} + \frac{dy_4}{dy_1} \frac{dy_1}{dp_{1g}} \\ &= (\beta_3 \beta_2 \beta_1 + \beta_3 \beta_1 + \beta_2 \beta_1 + \beta_1) \lambda = \pi_4 \end{aligned}$$

The total peer effect on the take-up of coworker j is given by π_j divided by the first stage coefficient λ . By comparing the estimated π 's across coworkers, we can identify the snowball effects. The second coworker identifies the direct effect, β_1 , as π_2 divided by λ . Subtracting off this direct effect, the snowball effect on the third coworker, $\beta_2 \beta_1$, is given by $\pi_3 - \pi_2$ divided by λ ; the snowball effect on the third coworker, $(\beta_3 \beta_2 \beta_1 + \beta_3 \beta_1 + \beta_2 \beta_1)$, is given by $\pi_4 - \pi_2$ divided by λ .

To estimate the snowball effects, we follow the RD design described in subsection 3.3. The reform generates quasi-random variation in the cost of taking leave for peer fathers whose children are born in a window surrounding the reform. We can estimate λ as the jump in take-up at the reform date cutoff in the first stage RD regression. We then estimate the π 's from a reduced form RD regression, examining whether this quasi-random variation in cost for the peer father changes the leave taking behavior of the coworkers. Using the differences in the estimated π 's, we estimate the snowball effects.

For simplicity of presentation, the system of equations above implicitly assumes the direct and indirect peer effects are independent of when the coworker's child is born.³² In reality, the influence of a peer is likely to decay over time, with a smaller peer effect for coworkers having children temporally distant from the peer father. This decay is a nuisance parameter which is not of immediate interest, but which must be accounted for in order to consistently estimate the snowball effects.

8.2 Empirical Results

Table 11 displays the estimated total reduced-form peer effect for each coworker in a firm. Note the first stage coefficient λ is the same for all coworkers ($\hat{\lambda}=.318$), and therefore does not affect the relative size of the snowball effect compared to either

³²A possible concern is that the spacing between a coworker's birth and the peer father's birth is affected by the reform. In Appendix Figure A1, we verify that spacing is continuous through the reform and therefore not a source of bias.

the direct or total peer effect. With this note in mind, we proceed by using estimates of the π 's to decompose the total peer effect into the direct effect and the snowball effect, and graph the relative importance of these effects over time.

The first column presents coworker estimates which do not account for decay. This regression mirrors our baseline specification, but allows for a separate discontinuity for each coworker. Even without subtracting out decay, the total reduced-form peer effect increases in magnitude from coworkers 2 through 4. Because we do not have enough observations to separately estimate effects within the firm for fifth and later coworkers, we estimate the average peer effect for this group. The total peer effect for this group declines, which is not surprising if decay is sizeable for this group who have children temporally distant from the peer father.

To identify the snowball effects, it is necessary to account for decay. The reason is that higher order coworkers have births which occur later in our sample period when there is more decay. To estimate decay, we exploit the fact that coworker 2 does not experience a snowball effect since there are no intermediate births in between him and the peer father (coworker 1). Hence, any change over time in the estimated peer effect for coworker 2 can be attributed to decay. We run a preliminary RD regression using the subsample of coworker 2 observations to estimate decay, and then adjust the estimates appearing in column 1 of Table 11 to account for depreciation.³³

To estimate decay, we augment equation (8) to include a polynomial in the timing difference between the birth date of the peer father's and coworker 2's child and an interaction term between these polynomial terms and the reform cutoff. The coefficients on the interaction terms divided by the coefficient on the reform cutoff identify the depreciation parameters. We plot the implied decay over time based on these estimates in Figure 9. Interestingly, the peer effect appreciates for the first 1.7 years before starting to decline again, with the depreciation term not becoming negative until approximately 3.5 years. This pattern makes sense once one realizes when fathers take leave from their firm. Although fathers generally sign up for leave before the birth of their child, most fathers do not begin their 4 weeks of leave until approximately 9 to 11 months after their child's birth³⁴. While there is likely to be some information transmission immediately after the birth of the peer father's child to coworker 2 (e.g., coworker 2 knows the peer father has already signed up for leave), more information is revealed after the peer father returns to work. At

³³Because we estimate the decay parameters based on a coworker 2 subsample, and because we do not have enough coworker 2's who have births later than 2002 in our dataset (less than 5%), our estimates of the snowball effects are restricted to 1993 to 2002.

³⁴This is because mothers and fathers cannot take leave at the same time, and mothers generally take all their leave before the father starts his leave. Many mothers exercise the option to take 48 weeks of leave at 80% earnings replacement rather than 38 weeks at 100% replacement.

that point, not only does the peer father have first-hand experience taking leave, but there is also an opportunity in the ensuing months to observe how the employer treats the peer father after his return to work.

Our key results are found in columns 3 through 6, which report decay-adjusted estimates. Whether decay is modeled as a third or fourth order polynomial makes little difference to the estimates. Focusing on the specification which allows for cubic decay, the reduced form coefficient for coworker 2 is estimated to be .028. This coefficient represents only the direct influence of the peer father on coworker 2, since there are no intermediate coworkers to create a snowball effect. It is smaller compared to the estimate in column 1, since on average, coworker 2's have their children early on when there is still appreciation. For coworker 3, the total reduced form peer effect rises to .038. The snowball effect accounts for 26 percent of the total peer effect.³⁵ As expected, the snowball effect is even larger for coworker 4 since there are more intervening coworkers. Coworker 4's reduced form estimate is rises to .064, with 56 percent of the total peer effect attributable to the snowball effect. For fifth and higher coworkers as a group, the snowball effect accounts for 66 percent of the total effect. As the table documents, these snowball effects are jointly statistically significant after accounting for depreciation. The pattern of increasing peer effects with each subsequent coworker captures the amplification of the original peer father's influence over time within a firm.

Figure 10 graphs the relative importance of the direct peer effect and the snowball effect over time, allowing both the direct peer and snowball effects to decay. The top line in the graph shows the actual leave take up for all coworkers having children after the original peer father. The bottom line subtracts the estimated total peer effect from the total leave take up. To construct this line, we use the estimated effects from column 3 in Table 11 to predict the size of the peer effect originating from the peer father (coworker 1), accounting for the mix of births in each year (coworker 2, 3, 4, and 5+ births) and adding back in depreciation.³⁶ The difference between the upper and lower lines illustrates how much lower the take up of leave would have been in each year had the original peer father not influenced any of his coworkers, either directly or indirectly. Figure 10 shows that this counterfactual gap, which includes depreciation, is sizeable and actually gets slightly larger over time.

Even more interesting is the decomposition of this counterfactual gap into direct

³⁵The percent of the total peer effect accounted for by the snowball effect for coworker j is calculated as $(\hat{\pi}_j - \hat{\pi}_2)/\hat{\pi}_j$.

³⁶We only plot the period from 1993 to 1999; extrapolating past 1999 is noisy and actually implies depreciation in excess of 100%. A likely reason is that most of our coworkers have their children long before 1999; the majority of observations past 1999 are to 5th or higher order coworkers in a firm.

peer effects and indirect snowball effects over time. The dark gray area in the graph indicates the direct effect of the peer father on coworkers, while the light gray area indicates the snowball effect. In 1993, virtually all of the estimated effect can be attributed to the direct effect, as there is little opportunity for intervening births to create a snowball effect. However, over time, the direct effect which can be mapped back to the original peer father (coworker 1) decays. By 1999 decay is large enough that the original peer father's direct effect on a coworker completely fades away. In contrast, the snowball effect gets larger and larger over time as more coworkers have a child within a given firm. Even though the snowball effects also decay, the accumulation of effects from intervening coworkers more than offsets this decay. By the end of the period, the snowball effect makes up almost 100 percent of the predicted total peer effect.

Figure 10 illustrates how important early peers are for future take up of social programs. From 1993 to 1999, program participation went from a little over 50 percent to over 70 percent of eligible coworkers. Much of this increase is due to common time effects, such as changes in societal norms, and the influence of other peer groups not captured by our estimates. However, even six years after the implementation of the program, the peer effects which can be traced back to the original father account for 21 percent of the total increase in program participation relative to 1993. This comparison highlights that peer effects can have long-lasting effects, even in the presence of decay, since any original peer effect cascades through a network over time. This is especially important for the rollout of new social programs, since our snowball findings indicate that participation rates early on can have long-lasting effects on future participation.

9 Conclusion

We find strong evidence for substantial peer effects of program participation in both workplace and family networks. Coworkers and brothers are 3.7 and 4.8 percentage points, respectively, more likely to take paternity leave if their peer father was eligible versus not eligible for paternity leave around the reform cutoff. These estimates imply sizeable peer effects of 11 and 15 percentage points for coworkers and brothers. The most likely mechanism is information transmission about costs and benefits, including increased knowledge of how an employer will react, and not leisure complementarities. We find substantial heterogeneity based on the strength of interpersonal ties between peers, with strong effects for long-term familial relationships and among male coworkers in small to medium sized firms, but no

evidence for peer effects in weaker extended workplace and family networks or in neighborhoods defined by geography. Finally, we find the estimated peer effect gets amplified over time within a firm, with each subsequent birth exhibiting a snowball effect in response to the original reform.

Taken together, our results have important implications for the peer effects literature and for the evaluation of social programs. Our study points out that individuals are affected by more than one peer group and that these social interaction effects are sizeable. The results illustrate that both the workplace and family can serve as important information transmission networks in settings where information is scarce and perceptions are in their formative stage. Our findings highlight that peer effects can have long-lasting effects, even in the presence of decay, since any original peer effect cascades through a network over time. This is especially important when considering the design and implementation of new social programs, since the initial group of participants can play a large and lasting role in the evolution of take up patterns. Social interactions can reinforce the direct effects on take up due to a program's parameters, leading to a long-run equilibrium take-up rate which can be substantially higher than in the absence of peer effects.

References

- AIZER, A., AND J. CURRIE (2004): “Networks or Neighborhoods? Correlations in the Use of Publicly-funded Maternity Care in California,” *Journal of Public Economics*, 88, 2573–2585.
- ANGELUCCI, M., AND G. DE GIORGI (2009): “Indirect Effects of an Aid Program: How Do Cash Transfers Affect Ineligibles’ Consumption?,” *American Economic Review*, 99, 486–508.
- ANGELUCCI, M., G. DE GIORGI, M. RANGEL, AND I. RASUL (2010): “Family Networks and School Enrolment: Evidence from a Randomized Social Experiment,” *Journal of Public Economics*, 94, 197–221.
- ATKINSON, A., L. RAINWATER, T. SMEEDING, O. FOR ECONOMIC CO-OPERATION, AND DEVELOPMENT (1995): *Income Distribution in OECD Countries: Evidence from the Luxembourg Income Study*. Organisation for Economic Co-operation and Development Paris.
- BABCOCK, P., K. BEDARD, G. CHARNESS, J. HARTMAN, AND H. ROYER (2011): “Letting Down the Team? Evidence of Social Effects of Team Incentives,” Discussion paper, NBER 16687.
- BAIRD, S., A. BOHREN, C. MCINTOSH, AND B. OZLER (2012): “Designing Experiments to Measure Spillover and Treshold Effects,” Discussion paper.
- BANDIERA, O., I. BARANKAY, AND I. RASUL (2009): “Social Connections and Incentives in the Workplace: Evidence from Personnel Data,” *Econometrica*, 77, 1047–1094.
- (2010): “Social Incentives in the Workplace,” *Review of Economic Studies*, 77(2), 417–458.
- BANDIERA, O., AND I. RASUL (2006): “Social Networks and Technology Adoption in Northern Mozambique,” *The Economic Journal*, 116, 869–902.
- BAYER, P., S. ROSS, AND G. TOPA (2008): “Place of Work and Place of Residence: Informal Hiring Networks and Labor Market Outcomes,” *Journal of Political Economy*, 116, 1150–1196.
- BERTRAND, M., E. LUTTMER, AND S. MULLAINATHAN (2000): “Network Effects and Welfare Cultures,” *Quarterly Journal of Economics*, 115, 1019–1055.

- BOBONIS, G., AND F. FINAN (2009): “Neighborhood Peer Effects in Secondary School Enrollment Decisions,” *Review of Economics and Statistics*, 91, 695–716.
- BURSZTYN, L., F. EDERER, B. FERMAN, AND M. YUCHTMAN (2012): “Understanding Peer Effects in Financial Decisions: Evidence from a Field Experiment,” Discussion paper.
- CARRELL, S., R. FULLERTON, AND J. WEST (2009): “Does your Cohort Matter? Measuring Peer Effects in College Achievement,” *Journal of Labor Economics*, 27, 439–464.
- CARRELL, S., AND M. HOEKSTRA (2010): “Externalities in the Classroom: How Children Exposed to Domestic Violence affect Everyone’s Kids,” *American Economic Journal: Applied Economics*, 2, 211–228.
- CARRELL, S., M. HOEKSTRA, AND J. WEST (2011): “Is Poor Fitness Contagious? Evidence from Randomly Assigned Friends,” *Journal of Public Economics*, 95, 657–663.
- CARRELL, S., F. MALMSTROM, AND J. WEST (2008): “Peer Effects in Academic Cheating,” *Journal of Human Resources*, 43, 173–207.
- CARRELL, S., B. SACERDOTE, AND J. WEST (2011): “From Natural Variation to Optimal Policy? the Lucas Critique Meets Peer Effects,” Discussion paper, NBER 16865.
- CASE, A., AND L. KATZ (1991): “The Company You Keep: The Effects of Family and Neighborhood on Disadvantaged Youths,” Discussion paper, NBER 3705.
- CULLEN, J., B. JACOB, AND S. LEVITT (2006): “The Effect of School Choice on Participants: Evidence from Randomized Lotteries,” *Econometrica*, 74, 1191–1230.
- DING, W., AND S. LEHRER (2007): “Do Peers Affect Student Achievement in China’s Secondary Schools?,” *Review of Economics and Statistics*, 89, 300–312.
- DUFLO, E., AND E. SAEZ (2003): “The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment,” *Quarterly Journal of Economics*, 118, 815–842.
- DUNCAN, G., J. BOISJOLY, M. KREMER, D. LEVY, AND J. ECCLES (2005): “Peer Effects in Drug Use and Sex Among College Students,” *Journal of Abnormal Child Psychology*, 33, 375–385.

- EVANS, W., W. OATES, AND R. SCHWAB (1992): “Measuring Peer Group Effects: A Study of Teenage Behavior,” *Journal of Political Economy*, 100, 966–991.
- FIGLIO, D. N., S. HAMERSMA, AND J. ROTH (2011): “Information Shocks and Social Networks,” Discussion paper, NBER 16930.
- GAVIRIA, A., AND S. RAPHAEL (2001): “School-based Peer Effects and Juvenile Behavior,” *Review of Economics and Statistics*, 83, 257–268.
- GLAESER, E., B. SACERDOTE, AND J. A. SCHEINKMAN (1996): “Crime and Social Interactions,” *Quarterly Journal of Economics*, 111, 507–548.
- GRANOVETTER, M. (1973): “The Strength of Weak Ties,” *American Journal of Sociology*, 78, 1360–1380.
- HANUSHEK, E., J. KAIN, J. MARKMAN, AND S. RIVKIN (2003): “Does Peer Ability Affect Student Achievement?,” *Journal of Applied Econometrics*, 18, 527–544.
- HENSVIK, L., AND P. NILSSON (2010): “Businesses, Buddies and Babies: Social Ties and Fertility at Work,” Discussion paper, IFAU-Institute for Labour Market Policy Evaluation.
- HOXBY, C. (2000): “The Effects of Class Size on Student Achievement: New Evidence from Population Variation,” *Quarterly Journal of Economics*, 115, 1239–1285.
- IMBENS, G., AND J. ANGRIST (1994): “Identification and Estimation of Local Average Treatment Effects,” *Econometrica*, 62, 467–475.
- IMBERMAN, S., A. KUGLER, AND B. SACERDOTE (2009): “Katrina’s Children: Evidence on the Structure of Peer Effects from Hurricane Evacuees,” Discussion paper, NBER 15291.
- JACOB, B. (2004): “Public Housing, Housing Vouchers, and Student Achievement: Evidence from Public Housing Demolitions in Chicago,” *American Economic Review*, 94, 233–258.
- KATZ, L., J. KLING, AND J. LIEBMAN (2001): “Moving to Opportunity in Boston: Early Results of a Randomized Mobility Experiment,” *Quarterly Journal of Economics*, 116, 607–654.
- KLING, J., J. LIEBMAN, AND L. KATZ (2007): “Experimental Analysis of Neighborhood Effects,” *Econometrica*, 75, 83–119.

- KLING, J., J. LUDWIG, AND L. KATZ (2005): “Neighborhood Effects on Crime for Female and Male Youth: Evidence from a Randomized Housing Voucher Experiment,” *Quarterly Journal of Economics*, 120, 87–130.
- KREMER, M., AND D. LEVY (2008): “Peer Effects and Alcohol Use among College Students,” *Journal of Economic Perspectives*, 22, 189–3A.
- KREMER, M., AND E. MIGUEL (2007): “The Illusion of Sustainability,” *Quarterly Journal of Economics*, 122, 1007–1065.
- KUHN, P., P. KOOREMAN, A. SOETEVENT, AND A. KAPTEYN (2008): “The Own and Social Effects of an Unexpected Income Shock: Evidence from the Dutch Postcode Lottery,” Discussion paper, NBER 14035.
- LALIVE, R., AND M. CATTANEO (2009): “Social Interactions and Schooling Decisions,” *Review of Economics and Statistics*, 91, 457–477.
- LEFGREN, L. (2004): “Educational Peer Effects and the Chicago Public Schools,” *Journal of Urban Economics*, 56, 169–191.
- LUDWIG, J., G. DUNCAN, AND P. HIRSCHFIELD (2001): “Urban Poverty and Juvenile Crime: Evidence from a Randomized Housing-mobility Experiment,” *Quarterly Journal of Economics*, 116, 655–679.
- LUDWIG, J., ET AL. (2011): “Neighborhoods, Obesity, and Diabetes, a Randomized Social Experiment,” *New England Journal of Medicine*, 365, 1509–1519.
- LUDWIG, J., J. LIEBMAN, J. KLING, G. DUNCAN, L. KATZ, R. KESSLER, AND L. SANBONMATSU (2008): “What Can We Learn About Neighborhood Effects from the Moving to Opportunity Experiment,” *American Journal of Sociology*, 114, 144–188.
- MANSKI, C. (1993): “Identification of Endogenous Social Effects: The Reflection Problem,” *Review of Economic Studies*, 60, 531–542.
- MARKUSSEN, S., AND K. RØED (2012): “Social Insurance Networks,” Discussion paper, IZA 6446.
- MARTIN, J., B. HAMILTON, P. SUTTON, S. VENTURA, T. MATHEWS, S. KIRMEYER, AND M. OSTERMAN (2009): “Births: Final Data for 2009,” *National Vital Statistics Reports*.

- MAS, A., AND E. MORETTI (2009): “Peers at Work,” *American Economic Review*, 99, 112–145.
- MAURIN, E., AND J. MOSCHION (2009): “The Social Multiplier and Labor Market Participation of Mothers,” *American Economic Journal. Applied Economics*, 1, 251–272.
- MOFFITT, R. (2001): “Policy Interventions, Low-level Equilibria, and Social Interactions,” in *Social Dynamics*, pp. 45–82. Cambridge: MIT Press.
- MONSTAD, K., C. PROPPER, AND K. SALVANES (2011): “Is Teenage Motherhood Contagious? Evidence from a Natural Experiment,” Discussion paper.
- MUNSHI, K. (2003): “Networks in the Modern Economy: Mexican Migrants in the US Labor Market,” *Quarterly Journal of Economics*, 118, 549–599.
- NANDA, R., AND J. SØRENSEN (2010): “Workplace Peers and Entrepreneurship,” *Management Science*, 56, 1116–1126.
- REGE, M., AND I. SOLLI (2010): “The Impact of Paternity Leave on Long-term Father Involvement,” Discussion paper.
- REGE, M., K. TELLE, AND M. VOTRUBA (2009): “Social Interaction Effects in Disability Pension Participation: Evidence from Plant Downsizing,” *UiS Working Papers in Economics and Finance*.
- ROSS, S. L. (2009): “Social Interaction within Cities: Neighborhood Environments and Peer Relationships,” *University of Connecticut Working Paper 31*.
- SACERDOTE, B. (2001): “Peer Effects with Random Assignment: Results for Dartmouth Roommates,” *Quarterly Journal of Economics*, 116, 681–704.
- STINEBRICKNER, R., AND T. STINEBRICKNER (2006): “What Can be Learned about Peer Effects using College Roommates? Evidence from New Survey Data and Students from Disadvantaged Backgrounds,” *Journal of public Economics*, 90, 1435–1454.
- ZIMMERMAN, D. (2003): “Peer Effects in Academic Outcomes: Evidence from a Natural Experiment,” *Review of Economics and Statistics*, 85, 9–23.




Figure 1. Paternity leave take up for all eligible fathers, 1992 - 2006.

Notes: Fraction of eligible fathers taking paternity leave based on the birth year of their child. Vertical line denotes the reform cutoff of April 1, 1993. There are two observations for 1993: one for the first quarter of the year (before the cutoff) and a second for the remainder of the year.




Figure 2. Coworker's and brother's fertility.

Notes: The top graph is for coworkers and the bottom graph is for brothers. Each observation is the average number of children born to coworkers/brothers in a bin, based on the birthdate of the peer father's child. The top graph uses one week bins, the bottom graph uses two week bins. The plotted local linear regression lines are based on daily, individual-level data. Dashed vertical lines denote the reform cutoff of April 1, 1993, which has been normalized to zero. See notes to Table 2.




Figure 3. Peer father's eligibility.

Notes: The top graph is for coworkers and the bottom graph is for brothers. Each observation is the average number of children born to coworkers/brothers in a bin, based on the birthdate of the peer father's child. The top graph uses one week bins, the bottom graph uses two week bins. The plotted local linear regression lines are based on daily, individual-level data. Dashed vertical lines denote the reform cutoff of April 1, 1993, which has been normalized to zero. See notes to Table 2.




Figure 4. Fraction of peer fathers taking leave.

Notes: The top graph is for coworkers and the bottom graph is for brothers. Each observation is the average number of peer fathers taking paternity leave in one-week bins (coworkers) or two-week bins (brothers), based on the birthdate of their child. Dashed vertical lines denote the reform cutoff of April 1, 1993, which has been normalized to zero. See notes to Table 2.




Figure 5. Coworker’s leave take up.

Notes: Each observation is the average number of coworkers taking paternity leave in a bin, based on the birthdate of the peer father’s child. The top graph uses one week bins, the bottom graph uses two week bins. Dashed vertical lines denote the reform cutoff of April 1, 1993, which has been normalized to zero. See notes to Table 2.




Figure 6. Brother's leave take up.

Notes: Each observation is the average number of brothers taking paternity leave in a bin, based on the birthdate of the peer father's child. The top graph uses two week bins, the bottom graph uses four week bins. Dashed vertical lines denote the reform cutoff of April 1, 1993, which has been normalized to zero. See notes to Table 2.




Figure 7. Local linear regression graphs for coworker’s and brother’s leave take up.

Notes: The plotted local linear regression lines are based on daily, individual-level data. The top graph is for coworkers and the bottom graph is for brothers. For comparison, dots for the average number of coworkers/brothers taking paternity leave in one week intervals (coworkers) and two week intervals (brothers) are also included in the figure, based on the birthdate of the peer father’s child. Dashed vertical lines denote the reform cutoff of April 1, 1993, which has been normalized to zero. See notes to Table 5.




Figure 8. Placebo estimates of the peer effect.

Notes: Each placebo estimate first assigns a window around a false reform date (i.e., a false cutoff), and then uses an RD to estimate a reduced form peer effect. There are 730 estimates for each graph (2 years of estimates), where each estimate increases the false reform date by one day. Placebo windows start after the true reform date to avoid being influenced by any jump at the true cutoff. The value of the estimate based on the true reform date (from Table 2) is labeled with a dashed vertical line in each graph.




Figure 9. Decay in the estimated peer effect over time.

Notes: Decay based on a sample of coworker 2's (i.e., the first coworker to have a birth after the peer father; these coworkers experience no snowball effect). To estimate decay, we augment equation (8) to include a polynomial in the timing difference between the birth date of coworker 2's and the peer father's child and an interaction term between these polynomial terms and the reform cutoff.




Figure 10. Relative importance of the direct peer and snowball effects over time.

Notes: The top line in the graph shows the actual leave take up for all coworkers having children after the original peer father. The bottom line subtracts the estimated total peer effect originating from the peer father from the actual leave take up using estimates from Table 11, accounting for the mix of births in each year (coworker 2, 3, 4, and 5+ births), and adding back in depreciation. The decomposition of the direct peer effect and the snowball effect is also based on estimates from Table 11.

Table 1. Timing of fertility around the reform window of April 1, 1993.

Birthdate of child	Coefficient
March 4 - 10, 1993	1.44 (4.58)
March 11 - 17, 1993	2.21 (4.58)
March 18 - 24, 1993	-3.05 (4.58)
March 25 - 31, 1993	-9.92** (4.58)
April 1 - 7, 1993 (first week post reform)	10.72** (4.58)
April 8-14, 1993	4.27 (4.58)
April 15-21, 1993	2.74 (4.58)
April 22-28, 1993	2.10 (4.58)
N	5,479

Notes: Regression of daily birth rates on dummy variables for birth weeks around the reform window. Control variables include day of week, month, and year dummies, as well as 365 day of year dummies. Sample includes all births between 1992 and 2006 to fathers eligible for any type of parental leave. On average, there are 840 births per week to eligible fathers in all of Norway. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2. Baseline regression discontinuity estimates for peer effects of fathers on their coworkers and their brothers.

	First stage (1)	Reduced form (ITT) (2)	Second stage (2SLS) (3)	N (4)
Panel A: Workplace network				
Fertility	.333*** (.024) [.03]	-.004 (.010) [.24]	-.011 (.029) [.24]	107,766
Take up of leave	.317*** (.026) [.03]	.035*** (.013) [.67]	.110*** (.043) [.67]	26,851
Panel B: Family network				
Fertility	.304*** (.011) [.027]	-.001 (.014) [.55]	-.002 (.047) [.55]	22,597
Take up of leave	.304*** (.014) [.026]	.047** (.020) [.57]	.153** (.065) [.57]	12,495

Notes: Both panels use daily data, exclude observations in a one-week windows on either side of the discontinuity, include linear trends in birth day on each side of the cutoff, and employ triangular weights. Control variables are father's and mother's years of education the year before birth, father's and mother's age and age squared at birth, parent's county of residence and marital status the year before birth, and an indicator for gender of the child. Panel A uses a one year window with 6 months on each side of the reform cutoff date, and restricts the sample to firms which have only one birth of any parity to eligible male employees in the one-year interval straddling the reform and coworkers whose first child is born after the peer father's child and the reform. Panel B uses a two year window with one year on each side of the reform cutoff date, and restricts the sample to fathers eligible to take any type of leave with a child of any parity born within one year of the reform and brothers whose first child is born after the peer father's child and after the reform. Column (1) is the RD estimate of the peer father's leave take up at the reform discontinuity, column (2) is the reduced form, or intent to treat, RD estimate of brother's fertility and leave take up based on peer father's eligibility at the reform discontinuity, and column (3) is the estimated peer effect. Standard errors clustered by firm in panel A and by family in panel B. Comparison mean in brackets based on coworkers or brothers of peer fathers with births in the pre-reform window. *p<0.10, **p<0.05, ***p<0.01.

Table 3. Specification checks for coworker and brother peer effects.

	First stage (1)	Reduced form (2)	Second stage (3)	N (4)
Panel A: Workplace network				
Baseline	.317*** (.026)	.034** (.013)	.109** (.043)	26,851
No Controls	.318*** (.024)	.034** (.012)	.106*** (.040)	26,851
Quadratic trend	.321*** (.041)	.043** (.021)	.134** (.068)	26,851
Cubic trend	.298*** (.062)	.050 (.032)	.168 (.111)	26,851
No donut	.323*** (.024)	.024* (.013)	.074* (.040)	27,856
Two week donut	.311*** (.028)	.042*** (.015)	.135*** (.050)	25,736
Non-eligibles included	.247*** (.021)	.033*** (.012)	.133*** (.049)	34,749
Cluster s.e.'s on day of birth	.317*** (.026)	.035*** (.013)	.110*** (.043)	26,851
Panel B: Family network				
Baseline	.304*** (.014)	.047** (.020)	.153** (.065)	12,495
No controls	.303*** (.013)	.046*** (.018)	.152*** (.059)	12,495
Quadratic trend	.319*** (.021)	.062** (.030)	.193** (.094)	12,495
Cubic trend	.329*** (.029)	.080* (.042)	.245** (.129)	12,495
No donut	.308*** (.013)	.043** (.019)	.141** (.061)	12,779
Two week donut	.303*** (.015)	.042** (.021)	.138** (.068)	12,204
Non-eligibles included	.220*** (.011)	.043*** (.017)	.197*** (.075)	17,835
Cluster s.e.'s on day of birth	.304*** (.014)	.047** (.020)	.153*** (.066)	12,495

Notes: Specifications mirror the baseline specifications described in Table 2. Standard errors clustered by firm in panel A and by family in panel B. *p<0.10, **p<0.05, ***p<0.01.

Table 4. Window robustness checks for coworker and brother peer effects.

Window	First stage (1)	Reduced form (2)	Second stage (3)	N (4)
Panel A: Workplace network				
90 days	.312*** (.036)	.043** (.018)	.138** (.060)	14,069
135 days	.320*** (.028)	.035** (.015)	.109** (.047)	20,498
180 days (baseline)	.317*** (.026)	.034** (.013)	.109** (.043)	26,851
Panel B: Family network				
180 days	.318*** (.020)	.063** (.029)	.198** (.091)	6,083
275 days	.309*** (.016)	.053** (.023)	.171** (.074)	9,179
365 days (baseline)	.304*** (.014)	.047** (.020)	.153** (.065)	12,495

Notes: Specifications mirror the baseline specifications described in Table 2, changing the window size on each side of the reform. Standard errors clustered by firm in panel A and by family in panel B. *p<0.10, **p<0.05, ***p<0.01.

Table 5. Local linear regression estimates for coworker and brother peer effects.

Bin width	First stage (1)	Reduced form (2)	Second stage (3)	N (4)
Panel A: Workplace network				
60 days	.317*** (.047)	.045* (.024)	.141* (.085)	9,030
90 days	.313*** (.037)	.042** (.018)	.134** (.063)	13,939
120 days	.306*** (.030)	.039** (.016)	.128** (.056)	18,055
Panel B: Family network				
120 days	.316*** (.025)	.066** (.033)	.208** (.104)	4,079
180 days	.312*** (.020)	.050* (.027)	.160* (.083)	6,052
240 days	.307*** (.017)	.052** (.023)	.170** (.071)	8,104

Notes: Samples mirror the baseline samples described in Table 2. Estimates based on local linear regressions with a uniform kernel with no control variables included. N is based on the number of observations in the bin width. Bootstrap standard errors, clustered by firm in panel A and by family in panel B, based on 2,000 replications in parentheses. *p<0.10, **p<0.05, ***p<0.01.

Table 6. Regression discontinuity estimates for direct effects of the reform on other outcomes.

	Total Years Employed (max=12)		Total Earnings (12 year annuity)		GPA of child	Married (in 2006)	# kids (in 2006)
Father	Mother	Ratio F/M	Father	Mother	Ratio F/M	(in 2006)	(in 2006)
(1)	(2)	(3)	(3)	(4)	(6)	(8)	(9)
-0.006	-0.038	-0.001	-6,261	-2,032	-0.000	.004	.005
(.031)	(.043)	(.002)	(3,793)	(1,636)	(.003)	(.007)	(.007)
[11.5]	[10.6]	[.47]	[356,707]	[195,871]	[.37]	[.67]	[2.6]
N 81,794	81,794	81,794	81,794	81,794	81,794	81,794	80,762

Notes: Specification uses daily data, includes linear trends in birth day on each side of the discontinuity, and employs triangular weights. Sample includes all fathers with a child of any parity born within one year of the reform who are eligible to take any type of parental leave. Standard errors clustered by extended family in parentheses. Comparison mean in brackets. *p<0.10, **p<0.05, ***p<0.01.

Table 7. Mechanisms in the workplace and family networks.

	First stage (1)	Reduced form (2)	Second stage (3)	N (4)
Panel A: Workplace network				
Peer father predicted to be a senior manager	.311*** (.049) [.023]	.072** (.031) [.66]	.233** (.103) [.66]	4,272
Peer father not predicted to be a senior manager	.316*** (.029) [.032]	.028* (.015) [.67]	.088* (.047) [.67]	22,579
Coworker in same firm when child is born	.335*** (.030) [.031]	.049** (.021) [.66]	.145** (.063) [.66]	10,576
Coworker not in same firm when child is born	.303*** (.029) [.030]	.025 (.017) [.67]	.085 (.055) [.67]	16,275
Peer father in private firm	.301*** (.027) [.028]	.051*** (.016) [.68]	.170*** (.055) [.68]	17,977
Peer father in public firm	.377*** (.041) [.038]	.032 (.029) [.74]	.084 (.077) [.74]	5,076
Peer father in industry with low unionization ($\leq 33\%$)	.358*** (.034) [.023]	.079*** (.026) [.69]	.219*** (.074) [.69]	6,834
Peer father in industry with high unionization ($> 33\%$)	.306*** (.028) [.033]	.036** (.017) [.70]	.117** (.055) [.70]	16,225
Peer father in firm with low tenure (< 10 years)	.307*** (.030) [.027]	.045*** (.016) [.66]	.148** (.053) [.66]	20,128
Peer father in firm with high tenure (≥ 10 years)	.328*** (.051) [.040]	.009 (.025) [.69]	.029 (.075) [.69]	6,723
Panel B: Family network				
Brother lives in same municipality	.304*** (.020) [.02]	.041 (.027) [.56]	.134 (.89) [.56]	6,673
Brother lives in different municipality	.307*** (.021) [.03]	.052* (.029) [.57]	.170* (.094) [.57]	5,723

Notes: Specifications mirror those in Table 2. Peer father is predicted to be a senior manager if he is the first or second highest earner in the firm. Coworker in same firm indicates the coworker is in the same firm when their child is born compared to 1993 (when the peer father had his child). Sample size can vary across subgroups due to missing values. Standard errors clustered by firm in panel A and by family in panel B. Comparison mean in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 8. Peer effects on additional days of leave.

	Reduced form (1)	Second stage (2)	Reduced form (3)	Second stage (4)
	A: Workplace network		B: Family network	
More than 1 week	.033** (.014) [.666]	.105** (.043) [.666]	.047** (.020) [.563]	.153** (.064) [.563]
4 weeks or more	.033** (.014) [.607]	.103** (.046) [.607]	.039** (.020) [.525]	.129** (.066) [.525]
More than 4 weeks	.009 (.011) [.172]	.028 (.035) [.172]	.032** (.014) [.148]	.104** (.047) [.148]
More than 8 weeks	.010 (.009) [.107]	.033 (.029) [.107]	.023** (.012) [.095]	.077** (.039) [.095]
More than 12 weeks	.008 (.008) [.083]	.025 (.025) [.083]	.016 (.011) [.074]	.051 (.034) [.074]

Notes: Specifications mirror the baseline specification described in Table 2 and have the same first stage estimates. N = 26,851 for the workplace network and N = 12,495 for the family network. Standard errors clustered by firm in the workplace network and by family in the family network. Comparison mean in brackets. *p<0.10, **p<0.05, ***p<0.01.

Table 9. Peer effects in networks with weak ties.

	First stage (1)	Reduced form (2)	Second stage (3)	N (4)
Panel A: Neighborhood network – two closest households on each side				
Fertility	.275*** (.009) [.03]	.001 (.005) [.22]	.004 (.017) [.22]	170,645
Take up of leave	.274*** (.012) [.03]	.002 (.012) [.58]	.008 (.043) [.58]	38,550
Panel B: Extended workplace network – husband of female coworker				
Fertility	.329*** (.034) [.043]	-.005 (.015) [.29]	-.016 (.046) [.29]	92,742
Take up of leave	.318*** (.037) [.040]	.015 (.016) [.52]	.047 (.049) [.52]	25,583
Panel C: Extended family network – husband of sister				
Fertility	.318*** (.013) [.044]	.011 (.017) [.54]	.036 (.052) [.54]	15,771
Take up of leave	.320*** (.017) [.043]	-.004 (.023) [.54]	-.013 (.072) [.54]	8,876

Notes: Specifications described in Section 7.1. Standard errors clustered by neighborhood in panel A, firm in panel B, and by family in panel C. Comparison mean in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 10. Workplace networks with strong versus weak ties.

	First stage (1)	Reduced form (2)	Second stage (3)	N (4)
Workplace network				
Start dates within one year of each other	.235*** (.032) [.023]	.059** (.027) [.64]	.251** (.120) [.64]	6,841
Start dates more than one year apart	.340*** (.029) [.033]	.024 (.015) [.68]	.071 (.045) [.68]	20,010
Firm size < 27	.313*** (.026) [.026]	.037* (.019) [.66]	.118* (.061) [.66]	13,229
Firm size ≥ 27	.312*** (.043) [.035]	.026 (.019) [.67]	.085 (.060) [.67]	13,662
Firm in rural areas	.294*** (.048) [.030]	.055* (.029) [.68]	.186* (.099) [.68]	6,099
Firm in urban areas	.322*** (.031) [.031]	.033** (.016) [.66]	.101** (.049) [.66]	20,173

Notes: Specifications mirror those in Table 2. Sample size can vary across subgroups due to missing values. Standard errors clustered by firm in parentheses. Comparison mean in brackets. *p<0.10, **p<0.05, ***p<0.01.

Table 11. Snowball effects on coworkers within a firm.

	Total effect (1)	Percent snowball (2)	Total effect (3)	Percent snowball (4)	Total effect (5)	Percent snowball (6)
	A: No decay		B: Cubic decay		C: Quintic decay	
Coworker 2 (π_2)	.034** (.017)	0%	.028** (.014)	0%	.027** (.013)	0%
Coworker 3 (π_3)	.037** (.017)	11 %	.038** (.017)	26%	.035** (.016)	25%
Coworker 4 (π_4)	.050*** (.018)	44 %	.064*** (.023)	56%	.060*** (.021)	55%
Coworker 5+ (π_5)	.025 (.016)	-18 %	.083 (.054)	66%	.078 (.051)	66 %
F-test for snowball p-value	1.01 [.387]		2.84 [.036]		2.84 [.037]	

Notes: Sample includes all coworkers having a child before 2002. Total effect is the total reduced form peer effect, accounting for decay as indicated in the specification headings. Snowball columns indicate the amount of the total effect that can be attributed to the snowball effect. The F-test for snowball effects is a joint test of $\pi_5 = \pi_4 = \pi_3 = \pi_2$. N = 22,869. *p<0.10, **p<0.05, ***p<0.01.




Figure A1. Spacing between the coworker's/brother's and the peer father's births.

Notes: The top graph is for coworkers and the bottom graph is for brothers. Each observation is the average number of days between births to a coworker/brother and the peer father in a bin. The top graph uses one week bins, the bottom graph uses two week bins. The plotted local linear regression lines are based on daily, individual-level data. Dashed vertical lines denotes the reform cutoff; the reform cutoff date of April 1, 1993 has been normalized to zero.

Table A1. Descriptive statistics for fathers in the workplace and family networks.

Father characteristics	One year window		Two year window	
	Coworker sample (1)	All fathers (2)	Brother sample (3)	All fathers (4)
Some college	.23 (.42)	.28 (.44)	.26 (.44)	.27 (.45)
Age at birth	31.3 (5.4)	31.9 (5.5)	28.9 (4.0)	31.9 (5.5)
Married	.45 (.50)	.48 (.50)	.39 (.49)	.48 (.50)
Child a girl	.50 (.50)	.49 (.50)	.49 (.50)	.49 (.50)
Number of children	2.7 (1.98)	2.8 (1.04)	2.7 (1.92)	2.8 (2.03)
N	7,504	38,958	10,823	81,913

Notes: Column (1) is our estimation sample of reform-window fathers in firms which have just one birth within 6 months on either side of the reform, and who also have a coworker whose first child is born after the father and after the reform. Column (2) is a comparison sample of all eligible fathers in Norway in the corresponding one year window. Column (3) is our estimation sample of reform-window fathers who have brothers, where the brother has a first child after the father and after the reform. Column (4) is a comparison sample of all eligible fathers in Norway in the corresponding two year window. There are 50, 134, 23, and 285 missing observations for the married variable and 166, 805, 68, and 1,684 missing observations for the some college variable in columns (1), (2), (3), and (4), respectively.

Table A2. RD estimates for direct effects of the April 1, 1993 reform on covariates.

	Workplace network (1)	Family network (2)
1. Father has some college	.034 (.030) [.22] 26,178	-.011 (.016) [.25] 12,340
2. Mother has some college	-.015 (.031) [.28] 26,502	.007 (.017) [.28] 12,240
3. Father's age at birth	-.375 (.371) [31.2] 26,851	-.106 (.164) [28.8] 12,495
4. Mother's age at birth	-.521 (.340) [28.7] 26,851	-.091 (.167) [27.1] 12,491
5. Marital status at birth	-.036 (.035) [.44] 26,708	.001 (.019) [.39] 12,495
6. Child is a girl	-.010 (.035) [.48] 26,427	-.001 (.019) [.49] 22,262
7. Father's firm size	-4.5 (5.0) [45.1] 26,851	– – – –
8. Father predicted to be eligible	.033 (.027) [.78] 34,385	.020 (.015) [.70] 17,696

Notes: Regressions use daily data, include linear trends in birth day on each side of the discontinuity, and employ triangular weights. Sample restrictions described in Table 2. For each regression, coefficient estimates, standard errors in parentheses, Standard errors in parentheses, clustered by firm in column (1) and by extended family in column (2). Comparison mean in brackets based on peer fathers with births in the pre-reform window. Number of observations reported below the comparison means. *p<0.10, **p<0.05, ***p<0.01.