Published in final edited form as:

Int Econ Rev (Philadelphia). 2016 August ; 57(3): 827–856. doi:10.1111/iere.12177.

Playing the Fertility Game at Work: An Equilibrium Model of Peer Effects

Federico Ciliberto,

Department of Economics, University of Virginia

Amalia R. Miller,

Department of Economics, University of Virginia

Helena Skyt Nielsen, and

Department of Economics and Business, Aarhus University

Marianne Simonsen

Department of Economics and Business, Aarhus University

Abstract

We study workplace peer effects in fertility decisions using a game theory model of strategic interactions among coworkers that allows for multiple equilibria. Using register-based data on fertile-aged women working in medium sized establishments in Denmark, we uncover negative average peer effects. Allowing for heterogeneous effects by worker type, we find that positive effects dominate across worker types defined by age or education. Negative effects dominate within age groups and among low-education types. Policy simulations show that these estimated effects make the distribution of *where* women work an important consideration, beyond simply *if* they work, in predicting population fertility.

Keywords

Career-family conflict; workplace interactions; multiple equilibria

1. Introduction

The demographic transition to lower fertility rates may have been a catalyst for sustained economic growth in previous centuries (Galor 2005), but sub-replacement fertility rates are now a major policy concern in much of the industrialized world. Fertility decisions affect the size and composition of the population. Reduced or delayed childbearing in the present leads

¹We thank Steve Durlauf, Leora Friedberg, Bryan Graham, Ed Olsen, John Pepper and participants at the AEA meetings, CIREQ Applied Economics Conference on Fertility and Child Development, Economic Demography Workshop at the Population Association of America meetings, Nordic Summer Institute, Aarhus University, Mount Holyoke College, Rice University, University of Houston, and University of Oklahoma for insightful comments and suggestions. We thank Søren Brøgger and Anders Tofthøj for excellent research assistance. This research was supported by the NIH National Institute of Child Health and Human Development (1R03HD061514), the Dean of the College of Arts and Sciences and the Vice President for Research and Graduate Studies at the University of Virginia, the Villum Foundation and the Danish Research Council (275-09-0020). Part of this work was completed while Miller worked at the RAND Corporation. The contents of this paper are solely the responsibility of the authors and do not necessarily represent the views of the NIH or other sponsors.

to population aging and higher ratios of retired to working populations. Hence, very low fertility rates present challenges for the financing of public and private pension schemes (Borsch-Supan 2000; Blake and Mayhew 2006), for redistribution under the welfare state (Rangel 2003), and for overall economic growth (Lindh and Malmberg 1999). Governments in Europe and elsewhere have sought to increase the size of the workforce by enacting policies that encourage fertility (through regulation or public spending, Grant et al. 2004; Mumford 2007) or that encourage women to enter the paid labor market. Because female labor market participation is generally associated with lower fertility, these two policies may be in conflict.

This paper studies peer effects in the fertility decisions of working women. Measuring workplace peer effects in fertility can improve our understanding of how female labor market participation affects population fertility rates. If peer effects are important, and either positive or negative, their presence implies that a new workplace factor should be introduced into fertility models: *where* women work may matter as much as *if* they work. In addition to their direct effects, peer effects in fertility may be important for mediating the impact of policy and other demographic changes on overall population fertility and fertility of different groups of women. Previous studies have discussed how positive peer effects can amplify or dampen fertility responses to changes in the policy or economic environment (e.g., Kohler 2001; Kravdal 2002; Bloom et al. 2008). An example of a dampening story is the argument in Moffitt (1998) and Murray (1993) that social stigma (associated with out-of-wedlock childbearing) reduced the short-term fertility responses to changes in fertility incentives in US welfare policies. An amplification story is in Montgomery and Casterline (1998), who argue that multiplier effects from social learning and social influences hastened demographic transitions, through the diffusion of contraceptive technologies.

Following the economics literature on fertility, we model childbearing as a rational choice, responsive to financial incentives (Becker 1960; Willis 1973). However, unlike most of that literature, which studies decisions at the level of individual women or couples,² we explore peer effects in an equilibrium framework for peer groups of women working at the same establishment. This emphasis aligns with the shift in economic demography to consider social influences on fertility decisions. Studies have found evidence of peer effects using geographic areas or neighborhoods (Bloom et al. 2008; South and Baumer 2000; Crane 1991), schools (Evans et al. 1992), ethnic or religious groups (Manski and Mayshar 2003), families (Kuziemko 2006) and networks of friends (Behrman et al. 2002; Bernardi et al. 2007) as their primary social unit. A common feature of these studies is their focus on social or informational factors leading to *positive* peer effects. In contrast, this paper studies interactions in fertility decisions among female coworkers at the same physical work establishment that could be positive *or* negative.

One motivation for studying the workplace is precisely this variation. Unlike the usual social effects that tend to increase similarity among friends or neighbors, the workplace setting

²See Hotz, Klerman and Willis (1997) for a survey of the literature on the economics of fertility in developed countries. More recent studies have revisited the relationship between household income and demand for children (e.g., Lovenheim and Mumford 2013) and the impact of redistributive policy on fertility (e.g., Kearney 2004).

contains a complex mix of social and economic interactions that can produce positive or negative peer effects in fertility. A second motivation is the dramatic increase over the last half-century in the share of prime-age women working for pay, which has made the workplace an increasingly important setting in which to study women's fertility decisions. Third, the workplace can be used to define peer groups of individuals who work for the same organization at the same physical location and are potentially exposed to one another for several hours a day. Although workplace peer groups are primarily structured around economic production goals, a new set of studies has found evidence of peer effects at the workplace involving various behaviors other than fertility, including productivity (Mas and Moretti 2009; Bandiera, Barankay, and Rasul 2010), charitable contributions (Carman 2003), retirement savings (Duflo and Saez 2003), and paternity-leave taking (Dahl, Løken and Mogstad 2014).³

We model the interactions at a workplace among women deciding whether to have children as a complete information, simultaneous, static, discrete choice ("entry") game and solve for the Nash equilibrium outcomes.⁴ In an equilibrium of the game, each woman is maximizing her utility by having a child or not, taking the actions of coworkers as given; this implies an inequality condition for each woman. The full set of inequality conditions at each workplace is a system of simultaneous discrete choice equations, which in equilibrium must all be satisfied.⁵ Methodologically, this is the key novelty of our approach, that we model decisions of all the agents in the reference group as simultaneous and as the outcome of equilibrium behavior. Thus, the peer effects are the outcome of equilibrium behavior of all members of the reference group.

We are interested in heterogeneous interaction effects for different agents. Heterogeneity across agents implies that the system of simultaneous fertility choice equations can have multiple equilibria both in the number and identity of the women having children. The presence of multiple equilibria complicates the estimation of the parameters of the model because it causes the likelihood function predicted by the model to sum to more than 1, making the model "incomplete" (Tamer, 2003).

The approach in this paper builds on the methods developed in Ciliberto and Tamer (2009) to incorporate multiple equilibria in estimation without imposing any rules for equilibrium selection in the regions of multiplicity (as is done, for example, in Cohen, Freeborn, McManus 2011; Rennhoff and Owens 2012; Krauth 2006; Card and Giuliano 2013).

³Dahl, Løken and Mogstad (2014) find peer effects both within families and workplaces.

⁴The assumption of complete information has the attractive features of capturing the reality that agents in the model likely have information about the payoff functions of their coworkers that is not available to econometricians and allows us to avoid the possibility of ex-post regret that could arise with incomplete information. In addition, in a static game of incomplete information, agents may experience ex-post regret, which we do not think as plausible in a long-run equilibrium. Also, the challenge of multiple equilibria that we address here would not necessarily be avoided in a static game with incomplete information (e.g., Berry and Reiss 2007, and Berry and Tamer 2006). Finally, Grieco (2014) rejects a model where agents only have incomplete information, while he cannot reject a model where agents have only complete information. This is because, as Grieco (2014, page 312) explains, "the model is at the height of its flexibility under the complete information assumption – multiple equilibria is most common and different equilibria allow the widest range of equilibrium entry probabilities."

⁵The use of the Nash equilibrium concept implies a self-consistency condition on outcomes at each workplace. The use of a self-consistency condition to address both individual and aggregate behavior in studies of social interactions is discussed in Brock and Durlauf (2001), who apply a notion of self-consistency based on a rational expectations rule in which average behavior conforms to expected behavior. By contrast, this paper uses a complete information setting and derives the self-consistency condition from the equilibrium condition of the fertility game.

Ciliberto and Tamer (2009) show that there is a methodological trade-off between point identification and assumptions on the nature of the heterogeneity across agents, but that it is possible to identify the sets of parameters (partial identification) such that the choice probabilities predicted by the econometric model are consistent with the empirical choice probabilities estimated from the data.

Our approach extends Ciliberto and Tamer (2009) along two dimensions. First, we allow for the number of interacting agents to vary across markets (workplaces); this provides an additional source of exogenous variation in the data that enables us, in principle, to point-identify the parameters for the effects of the control variables when there is only one woman in fertility age at the workplace. Second, we estimate the heterogeneous effects by *type* rather than for each individual agent because each woman is only observed in one workplace. Notice that this approach allows us to study the importance of heterogeneity across women and its implications on the equilibria of the fertility game. A complementary approach would be to consider a dynamic structural model of fertility decisions, for example to study the timing of the birth decisions, but that would need restrictive assumptions on the nature of heterogeneity across women to ensure the uniqueness of equilibria that one needs to estimate those types of models.⁶ A reduced form version of such dynamic behavior is carried out in Hensvik and Nilsson (2010), who use panel data to estimate a conditional linear probability model, where the dependent variable is the fertility decision and the explanatory variables are the timing of the co-workers' childbirths.

Because we explicitly incorporate an equilibrium concept in our full structural model, we can compute marginal effects that incorporate both direct effects (from changes in individual characteristics and contextual effects) and those mediated through changes in peer behavior. Our model can also be used to identify workplaces with multiple equilibria consistent with the observable variables and error terms. In these areas, small policy shifts may trigger large behavioral shifts or "phase changes" (Brock and Durlauf 2007) when many individuals change their behavior at once. The approach developed in this paper can be applied more broadly to other settings in which social interactions have strategic components.

We use data drawn from detailed administrative records on the population of Denmark. These records allow us to link individuals to their coworkers and family members to construct a cross-sectional database of individual and peer fertility outcomes and predictors for 2002 to 2005. We find positive interactions in fertility decisions among coworkers in individual Probit models with a variety of different controls. The positive estimates persist when we impose the Nash equilibrium self-consistency conditions in the structural model. However, in both Probit and structural models, when we introduce *establishment-level* (and firm-level) random effects to capture common shocks affecting all women at a workplace, the endogenous peer effects reverse in sign (as in Ciliberto and Tamer 2009). We also find important evidence of heterogeneous effects for different types of workers: positive effects dominate across worker types defined by age or education, but negative effects dominate within age groups and among low-education types. In the models with heterogeneous peer

⁶Adda, Dustmann and Stevens (2011) study the dynamic fertility choices of single individuals. Their approach considers the timing of the fertility decisions but *not in a strategic context*.

effects, we find that nearly half of all workplaces are in a region of multiplicity, with more than one Nash Equilibrium outcome. The peer effects that we estimate are meaningful in magnitude shifting fertility rates by more than 4.9 percentage points across all models, which suggests that the preferences of a woman's co-workers affect her fertility outcomes. Indeed, our policy simulations show potentially large changes in fertility rates from reallocating workers across establishments.

This paper is organized as follows. Section II discusses the model and identification. Section III describes the data and reduced form estimates. Empirical results from the full model are presented in Section IV and Section V contains policy simulations. Section VI concludes.

2. Econometric Model of the Fertility Game at Work

This section describes our game theoretic model of strategic interactions in fertility decisions among coworkers and then outlines our estimation approach.

2.1 Equilibrium in the Fertility Game

Our theoretical model of the fertility game builds on the individual fertility model in Jones, Schoonbroodt and Tertilt (2011) by adding the possibility of peer effects from co-workers. Agents in the model are women of childbearing age who decide whether or not to have a child. Agents aim to maximize their utility, which is defined broadly to encompass consumption of market goods and services, engagement in activities that are personally or socially rewarding, and pleasure from motherhood. The direct utility from motherhood may depend on the quality (behavior or achievements) of the child, which in turn may depend on parental investments in child human capital and the productivity of those investments.

After deciding on childbearing, each agent chooses her time allocation and consumption bundle, including investment in child quality for women with children, to maximize her utility, subject to her budget and time constraints (represented by t). We define the maximum utility that an agent without a child achieves as $u_{NK}(t, n)$ and the maximum utility achieved by an agent with a child as $u_K(t, q, n)$. Each of these functions depends on the exogenous variation in resources and time available to the woman, and the utility from motherhood also depends on the woman's preference for child quality (or her productivity in producing child quality, q).

The utility functions for both mothers and childless women also depend on n , the number of peers who have children in the period. The effects of peer fertility may flow through mainly social or economic channels. There may be positive social peer effects stemming from mimicry or a desire to conform to group norms (Bernheim 1994), where increasing the number of peers with children directly increases U_K (or decreases U_{NK}).⁷ Direct social effects on utility can be negative instead if agents want to be different from their peers,

⁷This channel would correspond to "social influence" in Montgomery and Casterline (1996). The "social learning" channel (where individuals learn from the past experiences of their peers and resolve uncertainty about their own future payoffs from different actions) is not included in the peer effects we measure in our static model of complete information. However, to the extent that these factors are present (for example, if a woman at the establishment was especially successful or unsuccessful at balancing work and family in 2001 or 2002), they will appear either as common shocks to all women at the same work establishment or (imperfectly) correlated errors. We account for both of these possibilities in estimation of our full model.

though these are less commonly studied (as they are less likely to occur with self-selected groups such as friends or neighbors that are usually used to define peers).

The economic channels for peer effects can also be positive or negative. Motherhood itself is associated with lower wages and wage growth (e.g., Waldfogel 1998, and Miller 2011, in the United States and Nielsen, Simonsen and Verner 2004, in Denmark). By reducing the negative signal to employers about the productivity of working mothers, coworker fertility can lead to positive financial spillovers for mothers. Alternatively, childbearing (and leave-taking) by coworkers can increase the incremental costs to the employer from hiring temporary replacements or rearranging work-flows, which would lead to negative spillovers (such as lower returning pay or an increased risk of job loss) for mothers. Finally, competition in internal labor markets (Lazear and Rosen 1981) may create positive spillovers to women who forgo childbearing while their coworkers have children if they find it easier to be promoted or otherwise advance professionally during their coworkers' absences, which could generate negative peer effects in fertility. The potential for negative peer effects is likely greatest in the Danish regime, which combines generous leave schemes with relatively late access to daycare; see details in Section III below.

We model these effects as a set of static games of complete information played once at each workplace and use the Nash Equilibrium solution concept. For each individual woman, the net utility gain from having a child can be represented as $v(t, q, n) = u_K(t, q, n) - u_{NK}(t, n)$, meaning that agents prefer to have a child if $v(t, q, n) \geq 0$. A set of fertility decisions (to have a child or not for each agent) is an equilibrium outcome of the game if no individual agent can improve her well-being by individually changing her action, taking the actions of all other agents at the workplace as given.

In the homogenous peer effects version of the model, the decisions of all other agents are summarized by a scalar n for each woman. In addition to this base case, we also consider cases with heterogeneous peer effects, in which women are grouped into discrete types and they potentially respond differently to the fertility decisions of other women of their same type and a different type (and this can also vary according to the woman's own type).

A simple example with two identical agents ($i = 1, 2$) is sufficient to illustrate some key features of the model. In this example, the economic problem can be summarized as:

$$\begin{aligned} d_1 &= 1 \text{ if } v(t, q, d_2) \geq 0 \\ d_2 &= 1 \text{ if } v(t, q, d_1) \geq 0 \end{aligned}$$

Each agent has a child if her net utility from childbearing is positive. In this example, it is clear that although the Nash Equilibrium conditions reject outcomes in which either woman would prefer to deviate, it does not ensure that the equilibrium outcome is efficient or that there is a unique solution for any particular game. For example, in the basic game with two identical agents and positive peer effects, there are two equilibria if $v(t, q, 0) < 0$, $v(t, q, 1) > 0$. Neither agent would want to deviate away from the equilibrium in which both agents have a child or from the one in which neither agent has a child.

Indeed, positive peer effects can generically lead to multiple equilibria in the fertility game with two or more agents. This means that for many workplaces, there may exist both lower-fertility and higher-fertility outcomes from which no individual woman wants to deviate. In the case of homogeneous interaction effects that are known to be negative, there may be multiple equilibria in the identity of agents who take the action of interest, but their number is uniquely determined (Bresnahan and Reiss 1990). However, this uniqueness is not guaranteed if the interactions are heterogeneous and there are more than two agents. For most of the workplaces in our study, this implies that multiple equilibria is a potential outcome that we need to address in estimation.

If the econometrician knew the selection function that determines the equilibrium being chosen by the agents, then the econometrician could exploit this information in estimation. However, the selection function is unknown, and it might depend on some or all of the observable variables, as well as on unobservable variables. For that reason, the selection function can be thought of as infinite dimensional nuisance parameter (Ciliberto and Tamer 2009). The econometric model is “incomplete” because the relationship between the unobservables, the exogenous variables, and the outcome is a correspondence and not a function (Tamer 2003). An important consequence for practical purposes is that the data would most likely reject any selection rule, even if were written as a function of market specific observable variables.⁸

Although our model is robust in the sense of accommodating a range of positive and negative peer effects and the resulting multiplicity of equilibria, we are only able to accomplish this by focusing on a static game. First, this means that we assume that the game is played only once by the agents. While it is true that women make fertility choices many times in their lives, this research will focus on a single cross-section of data at a particular point in time. We assume that this cross-section captures the long run equilibrium of fertility decisions within each establishment, which means that we would obtain the same empirical results, regardless of the particular time period we selected.⁹ We confirm the validity of this assumption in our data first by looking at fertility rates for women in our sample period year by year. These rates are quite stable over time, ranging from 0.079 to 0.081, reflecting the general stability of fertility rates for the Danish population during the period leading up to and including our sample period. The fertility rate in the overall female population aged 20–40 ranged from 0.064 to 0.069 between 1993 and 2005).

Second, our static game setup means that we do not model agents as playing repeatedly over time or responding dynamically to one another’s decisions and outcomes. This may be reasonable in the context of fertility choices within a couple of years, because neither researchers nor coworkers observe the time when the decision to have a child is taken. We observe instead the timing of births. Variation in the time between the decision and actual

⁸It is important to remark here that any test of a selection rule would presume that the model is correctly specified. For recent work on misspecification in models with moment inequalities, see Pomonareva and Tamer (2011).

⁹Or, as Schmalensee (1989, page 953) states, “the usual presumption in cross-section work in all fields of economics is that observed differences across observations reflect differences in long-run equilibrium positions.” This point is also in Berry and Reiss (2007). Static models of strategic interactions have a rich history in empirical industrial organization going back at least to Rosse (1970).

conception makes it impossible to determine the exact decision date.¹⁰ Thus, using data on births alone, we cannot determine which agent first decided to conceive.

Third, the game is played simultaneously. This means that all agents are assumed to make their fertility choices at the same time. In the case of fertility choices, this assumption seems particularly reasonable, because we do not observe the order with which agents made their fertility choices, and because the agents themselves are not immediately aware of their coworkers' decisions or pregnancies.

2.2 Empirical Specification of the Utility Function

The fertility game is played separately at each establishment, which we index by $e = 1, \dots, E$. The agents in the model are women of childbearing age. Agents are indexed by $i = 1, \dots, K_e$, where K_e is the number of female employees of childbearing ages at that establishment. In our complete model, we allow for heterogeneous peer effects by worker type. In that case, the utility to each agent from having a child depends on her own type and the fertility decisions of other agents at her workplace of each type.¹¹

An individual i at establishment e gains net utility of $v_{ie}(X_{ie}, n_e, \theta)$ from having a child. The variable n_e captures the peer effect experienced by the agent. In the single-type model, it is the count of other women at the establishment who have a child in the sample period. In the two-type extensions, it is a vector $n_{ie} = (n_{ie1}, n_{ie2})$, where n_{ie1} is the number of peers of type 1 who have a child in establishment e , and n_{ie2} is the number of peers of type 2 who have a child in establishment e . The variable N_e (or vector $N_e = (N_{1e}, N_{2e})$ in the model with two types) is the endogenous outcome of the game. Here, N_e includes all agents, while n_e includes only the peers of individual i .¹² There are $N_e + 1$ possible equilibria in the single-type model (the additional one is for the case when no woman at the establishment has a child), and $N_{1e} * N_{2e} + N_{1e} + N_{2e} + 1$ in the two-type game (the additional three cases are no woman has a child, only one woman of the first type has a child, and only one woman of the second type has a child). Let Γ denote the possible number of equilibria.

The vector X_{ie} consists of a set of observable exogenous variables that determine the net utility of an agent from having a child. These variables control for heterogeneity in observables across agents.

We estimate a linear approximation of this function as follows:

$$v_{ie} = \alpha X_{ie} + \sum_{r=1}^{R_e} \delta_r^{(i)} n_{re} + \varepsilon_{ie} \quad (1)$$

¹⁰According to Juul et al. (1999), 40% of fertile couples conceive within the first month, while 84% conceive within a year.

According to the same source, between 6 and 20% of European couples are infertile.

¹¹This treatment of heterogeneous effects differs from that in Ciliberto and Tamer (2009) where each agent's entry decision was allowed to have a different effect on the utility of each other agent, because the same sets of airlines were observed as potential entrants in multiple geographic markets. This is not feasible in our context of establishment peer effects, as each agent is observed as a potential child bearer only in one establishment.

¹²Notice that the problem is different from Ciliberto and Tamer (2009), where the outcome is a vector of binary values.

Establishments are indexed by $i = 1, \dots, K_e$ and worker types are indexed by r . Peer effects in fertility are captured by the $\delta_r^{r(i)}$ coefficients that relate increases in the fertility of type r peers on the net utility from childbearing for a type $r(i)$ agent. These terms measure peer effects within and between coworker types. In our empirical analysis, we allow $\delta_r^{r(i)}$ to be positive or negative, and will estimate its sign.

The error term ε_{ie} is the part of utility that is unobserved by the econometrician. We assume throughout that ε_{ie} is observed by all players in peer group e . Thus, this is a game of complete information. We write ε_{ie} as follows:

$$\varepsilon_{ie} = \varepsilon_{ie} + \varepsilon_e + \varepsilon_f,$$

where $\varepsilon_f \sim N(0, 1/2)$ is a firm-specific component; $\varepsilon_e \sim N(0, 1/2)$ is an establishment-specific component; and $\varepsilon_i \sim N(0, 1)$ is an individual-specific component. In some specifications we allow the individual-specific components to be correlated within establishments and we estimate the covariances.¹³ Examples of shocks that exhibit strong positive correlation at the establishment level are managerial attitudes and flexibility regarding working hours, promotions, and leave-taking, and the general degree of social and professional support available to help balance motherhood and work at the establishment. Notice that the establishment and firm specific shocks are each one drawn from normal distributions with variance $1/2$, so that their sum is equal to 1.¹⁴ By doing so, we assure that the individual woman shock has the same variance as the ones that are common among the women at the same establishment.

The outcome of the fertility game is defined in terms of types of agents rather than single agents. Thus, for a given value of the parameters, we will derive a predicted equilibrium outcome $N_e = (N_{1e}, \dots, N_{Re})$. Nevertheless, we will estimate the model and solve for equilibrium outcomes using individual utility functions.¹⁵

Three types of variables are included in the vector $X_{ie} = (Z_{ie}, S_e, W_{ie})$: individual characteristics (Z_{ie}), establishment specific characteristics (S_e), and instrumental variables (W_{ie}). Z_{ie} is a vector of individual characteristics that enter into the net utility of all the workers in the peer group, for example individual human capital that affects overall firm performance and wages for all workers. Elements of the Z_{ie} vector are included in estimation directly as factors that affect individual fertility; their establishment-level average values comprise the S_e vector of establishment level variables.¹⁶ W_{ie} contains individual characteristics that enter only into individual i 's utility. The values of W_{ie} for coworkers only influence the coworkers' utility from having a child, and do not directly influence the

¹³As discussed in Berry and Reiss (2006) and Berry and Tamer (2006) in the analogous context of unobservable fixed costs in oligopoly entry models, researchers are forced to impose distributional assumptions on the functional form of the shocks for estimation. Nevertheless, it is reassuring to note that our main results are robust to using a log-normal distribution for the error terms.

¹⁴Ciliberto and Tamer (2009) include three normally distributed shocks that are common to firms in a market. Each shock is drawn from a normal distribution with variance $1/3$, so the sum of the variances equals 1 in that paper as well.

¹⁵The fact that our sample includes workplaces with only 1 agent allows us to identify the parameters on the individual effects.

¹⁶We use median values for each variable over the set of agents in the work establishment. It is also possible to include establishment- or firm-level variables that are not based on elements of Z_{ie} in S_e . In our reduced form analysis, we found that these additional establishment variables had little effect on the estimated peer effects and we omit them from the structural model.

individual's own utility from having a child. This exclusion restriction helps us estimate separate endogenous and contextual fertility effects with greater precision, which is why we refer to the variables as instruments.

2.3 Estimation and Simulation

We estimate the parameters of the utility function in Section II.B with an approach that incorporates the possibility of multiple equilibria discussed in Section II.A by extending the econometric methods developed in Ciliberto and Tamer (2009) to estimate parameter bounds. The key strength of this approach is that it allows for heterogeneous peer effects that can take either sign and does not impose assumptions to select a single equilibrium. In this section we delve into the estimation procedure, explaining how the approach differs from Ciliberto and Tamer (2009).

Our unit of observation is an establishment. We separately consider establishments based on the numbers of women of fertile ages employed there, K_e . This is an important distinction from Ciliberto and Tamer (2009), where the number of potential entrants is the same in all markets. A crucial observation here is that 18 percent of establishments have only one fertile-aged woman. In these establishments, the woman's decision whether to have a child is not affected by coworker peer effects, and we exploit these observations to estimate the parameters that affect the fertility decisions directly (α) and not through the peer and contextual effects. The first set of parameters is estimated as if we were running a Probit, which makes their estimates very precise and stable. As a consequence, the estimates of the peer effects are also very precise.¹⁷

We simulate the distribution of the error term by drawing, for each woman, S values of each of the three components of the error term described above. Each draw is denoted by $s \in \{1, \dots, S\}$. As in Ciliberto and Tamer (2009), the draws are independent of the exogenous variables. This is akin to what is done in the random coefficients model (e.g., Berry, Levinsohn and Pakes, 1995), where the simulated tastes and demographic characteristics of consumers are drawn independently of product characteristics. If we instead drew errors that were correlated with the independent variables, we would artificially introduce endogenous relationships into the model (e.g., Train 2003; Berry and Tamer 2007, Result 4 in Section 2.4). The establishment and firm specific shocks are each one drawn from normal distributions with variance $\frac{1}{2}$, so that their sum is equal to 1.¹⁸ This sets the variance of the individual idiosyncratic shocks equal to the variance of the common shocks that apply to all women at the same establishment.

For each set of parameter values, we initialize two establishment-specific vectors of lower and upper bound counts of each possible equilibrium outcome, which we denote by H_e^L and H_e^U . The dimension of these vectors of probabilities is given by the total number of possible

¹⁷The estimates are remarkably robust to changes in the number of simulations. Our paper started with $S=40$, while Ciliberto and Tamer (2009) used $S=100$. We repeated our estimation with $S=100$ for the baseline specification and found very similar parameter estimates. In further robustness checks using simulated data, we found that parameter values were also stable when S was increased to 1000.

¹⁸In Ciliberto and Tamer (2009), there are three shocks that are common among firms in a market, and there each one is drawn from a normal distribution with variance $\frac{1}{3}$, so that their sum is equal to 1 there as well.

equilibrium outcomes in each establishment (in terms of number of births). At the start of the simulation routine, we set both H_e^L and H_e^U equal to zero and update the vectors at each simulation round.

We then compute the $\alpha_r X_{ie}$, to which we add a simulated draw of the unobservable, $\varepsilon_{ie}^{(s)}$; this sum is the value of net utility when peer effects are neglected:

$$v_{ie}(X_{ie}, n_{1e}, \dots, n_{Re}) = \alpha_r X_{ie} + \varepsilon_{ie} \quad (2)$$

We solve for the possible Nash Equilibria of the game by computing the hypothetical utility for each woman for any possible number of other women having a baby, here denoted by n_{re} , $r = 1, \dots, R_e$. For example, if there are two women of the same type r in a given establishment e , then this means that we have two possible utility values for each one of these women; one would be given by $\alpha_r X_{ie} + \varepsilon_{ie}^{(s)}$ when the other woman does not have a baby. The other one would be given by $\alpha_r X_{ie} + \delta_r^r + \varepsilon_{ie}^{(s)}$ when the other woman does have a baby.

The number of possible outcomes in establishment e is given by the product $\prod_{r=1}^{R_e} (1 + n_{re})$. As discussed above, the number of outcomes determines the cardinality of the vectors H_e^L and H_e^U . Thus, for example, in an establishment with two women, there are three possible outcomes: no women having children; one woman having a child; and both women having children. Notice that there are three possible outcomes because we only look at how many of the women are having a child, not which individual woman is doing so. Thus, in this establishment with two women, the two vectors initially are $H_e^L = (0, 0, 0)$ and $H_e^U = (0, 0, 0)$.

We determine all of the equilibria of the fertility game in each establishment for each simulation draw, s . For an outcome to be an equilibrium, it must be the case that no individual woman could be made better off by making a different choice. Every woman having a child in equilibrium must have a non-negative net utility from childbearing, conditional on the choices of all the other women. This condition ensures that no woman who has a child wants to deviate by not having a child. It must also be the case that no woman without a child would prefer to have a child. We test this by ensuring that the net utility of having children is negative for all agents who choose not to have children, conditional on the fertility of their coworkers.

Table 1 illustrates the possible equilibrium outcomes for the case of two women. The game has three potential outcomes, ranging from zero to two children. The outcome in which neither of the two women has a baby is an equilibrium as long as both $\alpha_r X_{1e} + \varepsilon_{1e}^{(s)} < 0$ and $\alpha_r X_{2e} + \varepsilon_{2e}^{(s)} < 0$, as shown in the first row. The outcome where only one woman has a baby is an equilibrium if either $\alpha_r X_{1e} + \varepsilon_{1e}^{(s)} \geq 0$ and $\alpha_r X_{2e} + \delta_r^r + \varepsilon_{2e}^{(s)} < 0$; or $\alpha_r X_{1e} + \delta_r^r + \varepsilon_{1e}^{(s)} < 0$ and

$\alpha_r X_{2e} + \varepsilon_{2e}^{(s)} \geq 0$. The analysis becomes much more complex with more agents and types of agents, as the number of conditions increases, but the principle is unchanged.

For every simulation draw s and establishment e , we count the number of equilibria consistent with the values of the observables and parameters. If there is a unique equilibrium, then we increment the corresponding entry in H_e^L and H_e^U by 1. If there are multiple equilibria, then H_e^L is unchanged and all the entries in H_e^U corresponding to the equilibria are increased by 1. In the example with two women, suppose that for the simulation $s = 1$ there are two equilibria, one in which no woman has a child and another in which both women have children (a possible outcome with relatively large and positive peer effects). In that case, $H_e^L = (0, 0, 0)$ and $H_e^U = (1, 0, 1)$. We repeat procedure for each the S draws (and all of the establishments) to obtain total counts. We use these counts to compute the lower and upper bounds on the probabilities of each potential outcome by dividing H_e^L and H_e^U by S .

Finally, we compare the predicted probabilities based on the rescaled values of H_e^L and H_e^U with the empirical probabilities (the rates at which we observe each of the possible outcomes) in the sample. Formally, the moment inequality conditions are given as follows:

$$H^L(\theta, X) \leq Pr(N|X) \leq H^U(\theta, X),$$

where θ is the set of all the parameters we wish to estimate, which includes both α and δ . $Pr(N|X)$ is the vector of empirical probabilities, which we estimate using a simple frequency estimator.

We apply the moment inequality conditions in estimation by minimizing the following distance function between empirical and predicted probabilities:

$$Q(\theta) = \int \left[\left\| \left(Pr(N|X) - H^L(\theta, X) \right)_- \right\| + \left\| \left(Pr(N|X) - H^U(\theta, X) \right)_+ \right\| \right] dF_x,$$

where $(A)_- = [a_1 1[a_1 - 0], \dots, a_T 1[a_T - 0]]$ for the lower bound. $(A)_+$ is defined analogously for $(A)_+$. This formulation means that a moment condition only adds to the distance function if the empirical probability is not included between the lower and upper bounds. The moment's contribution is a function of the distance, a , between the lower bound and the empirical probability if the latter is smaller than lower bound, or between the upper bound and the empirical probability if the latter is larger than the upper bound.

We use simulated annealing to find the parameters that minimize the distance. This method is far more robust than a grid search for functions such as ours that can be flat over ranges of parameter values.¹⁹ We record the optimized parameters from each iteration so that we can then map the whole distance function around its minimum value. The parameters associated to the values of the distance function around its minimum are then used to construct the confidence intervals.

We apply the methods in Chernozhukov, Hong, and Tamer (2007) to construct confidence regions that cover the identified parameter with a 95-percent probability. These methods are complex, as they require re-optimization of the distance function on subsamples of the original dataset and then comparing the minima for each of the subsamples to the values of the distance function in that subsample obtained using parameters from the full sample.²⁰ In practice, we only map the function around its minimum. We take the argument of each of these function values and compute the value of the distance function at these parameters in each of the subsamples. We then construct the differences between the values of the distance function at the parameter values that minimized the distance function on *that* subsample with the distance at the parameter values that minimized the distance on every *other* subsample. We repeat this exercise for each of the subsamples. For each subsample, we take the largest of these differences. We then sort these largest differences in increasing order of magnitude and choose the one that is at the 95th percentile. We take this value, divide it by the size of the original sample and sum it with the minimum value of the distance function in the original dataset. We obtain a critical function value that we use to construct the 95-percent confidence sets by collecting all the parameter values that make the function value lower than the critical value.²¹ For more information about this method, see the appendix in Ciliberto and Tamer (2009) and the original paper by Chernozhukov, Hong, and Tamer (2007). The computational burden of this approach is substantial but not unreasonable. If researchers have access to multiple processors and are able to run the code in parallel, they should be able to find a reliable minimum of the distance function within a week, even for large numbers of simulations.

2.4 Identification

The identification of social effects is challenging in any setting (see discussions in Manski 1993; Manski 1995; and Blume and Durlauf 2005). Manski (1993) identifies three reasons why individuals belonging to the same peer group may tend to behave similarly. First, individual behavior may be influenced by the behavior of other group members: endogenous effects. Second, individual behavior may respond to the exogenous characteristics of the group: contextual effects. These two comprise the social effects of interest. However, a third possibility is the presence of correlated effects in behavior that are unrelated to social interactions. This can occur if group members share similar observable or unobservable characteristics or face similar institutional environments. The fundamental problem of separately identifying these three effects from one another is denoted the reflection problem.

Even in the absence of correlated effects, the crucial novelty in our estimation framework is that we explicitly address the reflection problem by modeling the full system of individual equations. In the language of the peer effects literature, our parameter estimates separately

¹⁹In practice, the search is done starting from many initial values, including values for the exogenous variables that are found when we estimate the Probit regressions. We then vary the parameters of the simulated annealing in various ways: we consider different temperatures; different re-annealing intervals; different hybrid combinations of simulated annealing and other genetic algorithms and grid searches.

²⁰Each subsample is $\frac{1}{4}$ of the size of the original sample.

²¹In practice, we follow Ciliberto and Tamer (2009) and for each parameter we report the smallest and largest values that the parameter can take. That is, not every parameter in the “cube” belongs to the confidence region. This region can contain holes, but here we report the smallest connected cube that contains the confidence region.

capture endogenous effects and contextual effects. The endogenous effects are the direct effects of peer fertility on each individual's net utility from fertility (captured in the $\delta_r^{(i)}$ parameters on the peer fertility measures n_{re}). Contextual effects are estimated by including separate controls for individual values and establishment average values for elements of Z_{ie} . These are the individual factors that are assumed to predict both own and coworker fertility. In addition, because we estimate the variance-covariance matrix of the unobservables, we can also capture correlated effects in the unobservables. This is very different from Manski (1993), who studies the identification of a single equation model. It is also very different from most other empirical studies of social interactions that aim at identifying the sum of peer effects (endogenous and exogenous) or by assuming that only one type of effect is present.

In order to empirically distinguish endogenous fertility effects, the strategic effects of interest, from contextual effects that stem from the direct effect of coworker characteristics on one's own fertility, we employ an exclusion restriction.²² Specifically, we exploit the presence of a factor that affects the agent's own net utility function but is not likely to affect her peers' net utility functions (the variable W_{ie}), namely the agent's sisters' fertility. Recent evidence (Kuziemko 2006) indicates that women are more likely to have children after their sisters have recently had a child. While a woman's own sisters' fertility can increase her net utility from childbearing (because of social effects), there is no reason to expect the fertility of a woman's coworkers' sisters to have any effect on her net utility from childbearing (other than the effect that runs through peer fertility). For this strategy to work, we must therefore assume that having sisters does not directly affect a woman's fertility (this is supported empirically in footnote ²²) and that her family size has no direct effect on co-workers' fertility either.

We address the second potential challenge to identification from correlated errors within establishments in several stages, depending on the potential source of the correlation.²³ First, there may be random shocks at the establishment level or at the firm level that affect the utility from fertility for all women. The random utility shocks have a common firm specific, common establishment specific and an idiosyncratic individual component. Second, there may be correlations in the idiosyncratic shocks among women at the same establishment. In Section IV, we present estimates from models that allow for each of these types of random utility shocks.

3. Description of Administrative Data

We estimate the fertility game using data on the population of Denmark, a country in which women can exert a high degree of control over whether and when to have children. Abortion is legal and socially accepted (the ratio of abortions to births is about 1:4) and artificial insemination (AI) and in-vitro fertilization (IVF) are part of public health care that is

²²With large variation of excluded variables (Tamer 2003), one can get point identification of the latent payoff function. Here, however, the identification is only partial, as was the case in Ciliberto and Tamer (2009). Ciliberto and Tamer (2009, Theorem 2, p. 1802) show that exclusion restrictions can be used to reach point-identification or, heuristically, tighter sets of the parameter estimates.

²³Previous studies have addressed this problem by random assignment of individuals into peer groups (Sacerdote 2001) or by exploiting exogenous between-group variation (Graham and Hahn 2005). That source of variation is not available in our setting.

generally provided at no cost. The average age at first birth is 29 years and the average number of children per woman over the entire fertile age is 1.89. Roughly 12% of each cohort of women is voluntarily or involuntarily childless at age 49.²⁴

To understand women's fertility choices, it is important to know the institutional settings regarding parental leave, family-friendly policies, and childcare. The empirical analysis is conducted within a highly family friendly regime, especially aimed at and utilized by women. The Danish legislation offers generous parental leave in connection with childbirth and mothers can currently be on job-protected leave for up to 46 weeks after giving birth. The vast majority of mothers take leave that exceeds six months. Compensation during leave is typically higher for public sector employees than for women working in the private sector.

Daycare is, for the major part, publicly provided and organized within the municipalities and heavily subsidized. All children are eligible for publicly provided childcare and though some waiting time is common, slots are typically available when the child is around one year old. In general, opening hours during weekdays are between 6:30 am and 5:00 pm, in principle allowing both parents to work full time.

Our dataset uses several administrative registers maintained by Statistics Denmark. The primary data source is a merged employer-employee data set, which includes information on the entire population of Danes aged 15–70. The data set covers the period from 1980 to 2005. Using unique person and workplace establishment identifiers we link all coworkers at the establishment level. The data contain yearly information about socioeconomic variables such as gender, age, family status as well as family identifiers, education, labor market experience, tenure at current job, unemployment levels, leave-taking, and income. Of particular interest for our study, births are identified using exact birth dates from the national fertility register.

From this population we select the group of fertile-aged women (20 to 40) who worked at a given establishment within a firm in November 2002. Our sampling does not condition on employment status after November 2002, because this is endogenous and could bias our estimates. In our main analysis we focus on establishments with 10–50 employees thus excluding small establishments where it would be difficult to credibly distinguish peer effects from selection and large establishments where it would be difficult to identify groups of interacting coworkers. We observe 31,725 medium-sized establishments. Roughly 25% of employed women in the relevant age range worked in such medium-sized establishments in November 2002 (151,494 out of 647,678 employed women). Women working at medium-sized are similar to those working at larger establishments in the values of the control variables, though they are more likely to be in the private sector; women at small establishments stand out in comparison as less educated and even more likely to be in the private sector. Average fertility rates are lower at smaller establishments and higher and larger ones.

²⁴Source: <http://www.statistikbanken.dk>

We define coworkers as individuals at the same establishment based on their firm and location of work, as of November 2002. Our fertility outcome is an indicator variable for having a child between December 2002 and November 2005.

3.1 Descriptive Statistics

Panel A of Table 2 presents descriptive statistics for the individual level control variables used in our main estimates for the women in our sample, both overall and separately by fertility outcome. About 24 percent of the women in our sample gave birth in 2003 to 2005. The average age in our sample is about 30 years (mean age is 30.04 years, not reported in the table), and women above that age are more likely to have a child during the sample than those below it. Women with some college education are 30 percent of the full sample, but 40 percent of the sample of women with births. Like older women, women with more (above the mean of 6 years) work experience are more likely to have children.

On average, about half of the women in our sample were mothers prior to 2003. Previous fertility is a strong predictor of fertility during the sample period. Women with one child were far more likely to have a second during the sample period (they comprise over one-third of those giving birth during the sample period and less than 13 percent of those not giving birth), while those with 2 children were far less likely (comprising over 28 percent of the sample without births and less than 13 percent of the sample with births). This is consistent with a dominant two-child norm in Denmark and with birth spacing generally falling within three years. Women who gave birth in 2003 to 2005 were also, on average, younger and less educated. Finally, women who had a child in the sample period are significantly more likely to have a sister who previously had a child (2.3 percentage points). This is similar to the finding in Kuziemko (2006); it provides exogenous variation that we exploit for identification.²⁵

Panel B of Table 2 presents descriptive statistics at the establishment level. The mean number of employees is about 20 but over 60 percent of establishments employed fewer than that number of workers. Focusing on our peer group of interest, the average establishment employed roughly five women of reproductive age (20 to 40). Nearly 20 percent of establishments employed only one woman in the relevant age range. These women are not susceptible to workplace peer fertility effects, and they provide variation that is used to identify the effects of individual level explanatory variables. Although they are slightly different from other women (being younger and less likely to already have a child at the start of the sample period), the group includes substantial shares of women with each of the possible values for each of the variables in the model, which means that extrapolation beyond the range of values observed for these women is not necessary for estimation. Figure 1 plots the full distribution of peer group size (number of female workers of fertile ages) across establishments in our sample. Although our selection rule could accommodate up to 50 fertile-aged women per establishment, in practice we observe no more than 38. Figure 2 plots the distribution of our main fertility outcome: the number of women at each

²⁵Not reported in the table, we also found a highly significant gap in fertility rates between women with sisters who had children recently (0.272) and women without such sisters (0.232). In contrast, we did not find substantial differences in fertility rates (or individual controls) between women with and without sisters.

establishment who have a child. Figure 3 shows the distribution of the ratio of these variables: the share of women with births.

3.2 Individual Probit Estimates

Before investigating the results from our main model that incorporates the strategic interactions in fertility decisions, we explore the key conditional associations in the data using a series of Probit models for the individual propensity to have a child. We estimate the model with different sets of covariates to explore the sensitivity of the estimated associations between own and peer fertility to the inclusion of different controls. Unlike the structural estimates in Section IV, this analysis is performed at the individual level.

Table 3 reports the marginal effect and standard error estimates for the peer fertility effects under different Probit models. In the basic model with only peer fertility and a constant term, the relationship is positive and significant. Each additional coworker that has a child is associated with an increase in the average fertility propensity by about 1 percentage point, which is about 4 percent of the sample mean (of 23.7 percent, see Table 2).²⁶ The inclusion of individual characteristics as controls decreases the coefficient estimate to 0.3 percentage points (still highly statistically significant), but the estimated effect remains relatively stable as other variables are added to the conditioning set. For computational feasibility, we limit the set of controls to a parsimonious set of key “basic” controls for individual characteristics (indicators for age, education, and experience categories, as well as previous own and sisters’ fertility) and to a richer set that add controls for contextual effects (median age, education, experience and past fertility of peers). The marginal peer effect estimates from these models are also positive and highly statistically significant, at about 0.5 percentage points.

We also explored the robustness of the Probit estimates to changing the sample definition (results are available upon request). The estimates are not sensitive to considering births in 2003–2004 (instead of 2003–2005) or births in 2003 only. The smaller samples have less precise estimates that are qualitatively unchanged. When we restrict the analysis to women who stay employed at the same establishment in 2003–2004 or 2003–2005, the point estimates are slightly smaller, though there are no significant differences. These last estimates are difficult to interpret because leaving a company is an endogenous outcome that may be related to career experiences, including own or co-worker fertility. Therefore, in the main analysis, we define the peer group in November 2002, right before the sample period for fertility outcomes.

The one change to the model that had a qualitative effect on the peer effect estimate is the addition of establishment-specific random effects. Matching the structure we use in the full model below, we assume each of these shocks is normally distributed with a mean of zero and variance of 1. These shocks account for spurious positive correlations in fertility among coworkers that are related to workplace factors other than peer fertility (common shocks in

²⁶The elasticity is similar in magnitude to the marginal effect, because the probability that an individual has a baby is close to the share of colleagues who have a baby. For example when the estimated marginal effect is 0.01, the elasticity is equal to $(0.01/0.229) / (1/0.237) \approx 0.01$.

Section II.D). Their inclusion shifts the peer estimate to negative and statistically significant, as reported in the bottom panel of Table 3. This pattern is repeated in the full model below.

4. Simultaneous Estimation of the Fertility Game

This section presents estimates from the full structural model for the fertility game using the approach described in Section II. The unit of analysis is a work establishment and the outcome of interest is the total number of women having a child in the sample period (2003 to 2005). As described in Section II, our structural approach incorporates the Nash Equilibrium constraints requiring that each agent is enacting her own best response to her coworkers' fertility decisions, but does not make assumptions about which equilibrium will occur in cases of multiplicity. Hence, our peer effect results are bounds on the parameters rather than unique point estimates.

4.1 Homogeneous Peer Effects

Table 4 reports these parameter bounds for our initial structural specification for a single homogenous peer effect for all coworkers. Column 1 reports the basic model with individual-specific control variables. Column 2 uses the same set of control variables but allows for common firm-level and establishment-level random effects in the overall error term. Including these effects has a dramatic change on the estimated peer effects.²⁷ The bounds on the estimates change from positive in the model that assumes independent errors to negative when common shocks are included. This reversal of the sign of the basic relationship in the model with the inclusion of random effects highlights the importance of firm-level and establishment-level unobserved heterogeneity in determining fertility outcomes at different workplaces.

Although the inclusion of the firm-level and establishment-level errors has a major effect on the peer effect estimates, it does not change the qualitative estimates for the controls. In both models presented in the first two columns of Table 4, women over the age of 30 are more likely to have a child in the sample period, as are women who are college-educated, those with more years of experience and with a single child born before 2003. Mothers with two children by 2003 are significantly less likely to have a child in the period, consistent with the two-child norm mentioned in Section III.A. Column 3 adds the variable for sisters' fertility (an indicator for having a sister with a new child from 2000 to 2002) to the model; it has a positive fertility effect.

In the first 3 columns, all of the individual variables are included in the agent's own utility equation, but not in the net utility functions of the agent's coworkers, which means that the model effectively treats all these variables as instruments for own fertility. The basic model

²⁷We exclude income and tenure from our model because they are both related to the woman's current place of employment and may be endogenous to current and anticipated fertility. In a robustness check, we found that including the variables in our baseline model (Table 4, Column 2) leaves the estimated peer effects essentially unchanged. This is possibly because our basic conditioning set already absorbs the effects of income and tenure. Ideally one would want to account for the endogeneity of income and tenure. This is, however, challenging for two reasons. Firstly, it has proven difficult to find variation explaining income and tenure but not fertility outcomes (exclusion restrictions). Secondly, it is very difficult to correct for endogeneity of variables in highly non-linear models. A popular solution to correct for endogeneity in non-linear models is to use control functions, but there is no theoretical work that shows how that approach would work in the complex non-linear simultaneous equations model we consider here.

excludes all contextual effects; agents only affect their coworkers through the endogenous fertility decisions in the game. Although these may be reasonable exclusions, we relax them in the final column, where establishment-level median values of all of the variables except the designated instrument (sister fertility) are included in the model. Several of these contextual effects are found to be relevant predictors of fertility, but this expansion of the model leaves the estimated peer effects largely unchanged.

The parameter bounds in Table 4 reflect coefficients in the net utility function and there is no natural, direct interpretation for their magnitudes. Instead, we determine the quantitative importance of the effects by computing marginal effects from one-unit changes in each of the variables on average fertility rates for agents in the sample. These values are computed by comparing the predicted equilibrium fertility rates in the sample with each of the binary variables set to one or zero, and with the discrete experience measure set to its actual value or incremented by 1 unit (corresponding to 10 years).²⁸ The equilibrium fertility rates with higher and lower values for each control are shown in the first 2 columns of Table 5 and the marginal effect, the difference between these rates, is in the third columns. The top panel of the table reports these marginal effects with the peer effects assigned at their estimated parameter value and the bottom panel makes the same comparison but sets the peer effects to zero.

Whether or not we allow for endogenous peer effects in fertility, we find marginal effects of college education, work experience, age over 30, having one child before the sample period, and having a sister with a child between 2000 and 2002 that are all positive, while the effect of having two children before the sample period is negative. The largest effects are associated with age over 30 (8.54 and 9.59 percentage point increases, with peer effects on or off, respectively) and from having a previous child (8.98 and 10.01 percentage point increases). Turning on the indicator for sisters' fertility has the effect of increasing predicted fertility by about 1 percentage point (0.85 percentage points with peer effects and 0.96 percentage points without them).

The average marginal effect of a unit increase in peer fertility in our sample can be seen by comparing the predicted fertility values for experience equals zero in the top and bottom panels (with and without peer effects). Turning on peer effects leads to a 5-percentage point decline in the share of agents who are predicted to have a child in the sample period. This is a sizeable effect that represents over 20 percent of the average fertility rate in our sample (24 percent). Furthermore, the marginal effect of each of the control variables is larger when we ignore peer effects in fertility (in the bottom panel). This indicates that the negative peer effects lower fertility both directly, through women's best responses to their coworkers' fertility decisions, and indirectly, by moderating the impacts of other factors.

The finding of negative peer effects in fertility is consistent with several of the channels for peer effects described in the theoretical presentation in Section II.A. Specifically, although a desire to adhere to group norms or scale economies from information sharing and childcare

²⁸The parameter estimates in Table 4 are reported as bounds. In order to compute a single fertility rate outcome, we use the parameter values at which the distance function was minimized in estimation.

coordination would yield positive peer effects, negative effects could come from a desire to distinguish oneself from one's peers, from scale diseconomies when multiple workers take leave at once, or from tournament competition in internal labor markets. The overall negative effects are also consistent with some of the positive channels operating, but the results indicate that the negative channels dominate in the overall population.

One concern about these results is that the finding of negative average peer effects may in part be an artifact of our modeling decision (in Equation 1) for effects that increase linearly in the number of coworkers having a child. This model is appropriate if the incremental effect of each additional coworker having a child is constant and does not depend on the number of agents in the workplace. Appendix Table A1 reports estimates from an alternative version of the model, in which the peer effects enter the net utility function as a share rather than a count variable. This functional relationship is analogous to the usual peer effects measures in linear-in-means model of social interactions (Graham and Hahn 2005) where the average peer characteristic is what matters. The results are robust to this alternative measure of peer fertility. The model in column 1 with independent errors produces positive peer effects, but adding the random effects at the firm and work establishment in column 2 lead to negative peer effects.

Another concern that often arises in peer effects research is the potential for endogenous peer selection, whereby individuals are more likely to be peers with others who resemble them. In our workplace setting, the specific concern would be that women choose workplaces in part based on their fertility preferences (or an unobserved factor that affects fertility decisions), which could happen if establishments vary in the degree to which they provide a "family-friendly" environment that supports working mothers. Although maternity leave and childcare benefits are relatively generous for all Danish workers, public sector employees receive longer periods of paid leave with greater compensation, on average, and there is more variation across workplaces in the benefits available to private sector workers (based on their individual or collectively-bargained contracts). Both sectors likely contain some variation in the attitudes of coworkers and employers toward working mothers. We examine the importance of sorting by estimating our model separately for establishments in the public and private sectors.

We expect that sorting will produce a positive bias in the estimated peer effects and that this effect would be more severe in the private sector than the public sector and limited to models that do not include common error terms for women working at the same establishment and firm. This is what we find. When we estimate the model without these random effects separately by sector, we only find significant positive peer effects in the private sector, where sorting likely plays a larger role (Appendix Table A2). When we estimate the model with establishment- and firm-specific random effects, we find little sector difference. The peer effects are negative and significant for each of the sectors and the confidence bounds are overlapping. This indicates that the random effects are absorbing the common shocks that vary across establishments and increase the fertility of all coworkers (and potentially attract female workers with similar high tastes for fertility).

Finally, we present some evidence on the fit of the model to the data. Following the standard approach that is used for discrete choice models such as the Probit, we report the percentage of outcomes that are correctly predicted by our model. Notice that in each establishment the data only contains one given number of women having a child. The model, however, can predict multiple equilibria in the number of women having children. If one of predicted equilibria is the outcome observed in the data, then we conclude that our model predicted the outcome correctly. We find that our model predicts approximately 13 percent of the outcomes in the data. We can use this quantity to compare the fit of different models.

The main finding in this section of *negative* workplace peer effects in fertility is novel, but the analysis is limited by its reliance on the assumption that all agents have the same influence on all other agents at their workplace. We relax the assumption of homogeneous peer effects in the next section by estimating heterogeneous peer effects for different types of agents.

4.2 Heterogeneous Peer Effects by Worker Type

The different theoretical channels for peer effects in fertility, both social and economic, suggest that our model of homogeneous effect in the previous section may be too limiting. Social effects related to group norms might be expected to be stronger within groups if agents care more about being similar to their closer peers, leading to overall effects that are more positive within subgroups rather than across them. Alternatively, if the subgroups differ in social status, the social effects could lead all agents to want to resemble one group but not the other. In that case, members of the lower status group would prefer to imitate the higher status peers, leading to positive peer effects across subgroups and negative effects within subgroup for agents in lower status subgroup. Scale economies from sharing information could similarly be stronger within subgroups if women are more likely to coordinate with their more similar peers, or stronger across groups if sharing is more likely across type. The career competition effects could lead to more negative peer effects within subgroup (the set of closer competitors), particularly for subgroups engaged in more internal competition. The scale diseconomies for replacing a worker (and keeping their job open) could similarly be stronger (more negative) within subgroups that are related to job functions or tasks, and if these are large enough, workers who take leave at the same time as many similar coworkers may not have the option of returning to the same job at that employer after their maternity leave. This last effect can also differ across subgroups, in this case, depending on the scarcity of their skills and the costs to the firm of replacing them.

In this section, we estimate expanded versions of our model that allow for heterogeneous peer effects by binary worker types defined based on education (defined by college education versus less than college) and then by age (over or under age 30). For each version of the model, we estimate four peer effect parameters: the effects of the low-type on the low-type, of the low-type on the high-type, of the high-type on the low-type and of the high-type on the high-type. Workplaces with multiple agents of a single type are used to identify the within-type parameters and those with agents of different types provide identification for the cross-type parameters.

The model with separate effects by education type in Table 6 reveals important heterogeneity in the peer effects across different subgroups of women. The first three columns of the table show 95-percent confidence parameter bounds for the peer effects (by type) and control variables in expanded versions of the specifications in Table 4, columns 2 to 4. The final column of Table 6 adds a specification that allows for correlated errors across agents within each workplace and estimates an additional parameter for that correlation. Across all specifications, the pattern of the peer effects is unchanged. The overall negative peer effects from the homogeneous model in Table 4 are reflected in the negative peer effects among lower education women; these women comprise over two-thirds of the sample, so it may not be surprising that their effects dominate the overall average. However, the peer effects among higher education women are found to be *positive*, as are both of the cross-type peer effects. Recall that all models in Table 6 include random effects for firms and establishments and column 4 also allows for correlated individual-specific errors. The inclusion of these terms provides some protection against a spurious finding of positive peer effects (as in the first column of Table 4), which makes the positive effects more credible when they occur.

What do these results indicate about the operative channels for peer effects? By uncovering the importance of heterogeneous effects, and the specific finding that peer effects are positive for some groups and negative for others, we show that multiple channels, operating in opposing directions, are at play. Our stylized representation of the net utility function (with a limited set of exogenous controls) does not allow us to definitively pinpoint the exact mechanisms leading to these peer interactions. However, the results provide suggestive evidence regarding the potential channels at play. The negative peer effects for low-education women from the fertility of their low-education peers, but positive effects from the fertility of the higher-educated peers could indicate a social desire to mimic the higher-status, higher-education group. However, this pattern is not (exactly) what we find among higher-education women, who respond positively to fertility of both higher-educated and lower-educated peers. Similarly, the possibility of information spillovers being more valuable when the source is higher-educated women is also not consistent with the positive effects of fertility from low-education to high-education agents that we find.

Instead, we speculate that the pattern suggests that the social forces are positive both within and across subgroups, but that the negative effects of competition and scale diseconomies are relatively more important for low-education women. The fact that fertility of high-education women generates a smaller peer effect on other high-education women than the cross-type spillover effect to low-education women indicates that, even for high-education women, competition is dampening the positive within-type peer effect. However, it is not sufficient to turn the overall effect negative. Although competition for promotion (and there the incentive to delay childbearing when peers are having children) can be important for both education groups, it is possible that sex segregation in tasks is more prevalent for lower-education women (occupational sex segregation is decreasing in education level, Cotter, Hermesen and Vanneman 2004), which makes female coworkers a more important group of workplace competitors than for higher-education women who are more likely to also be competing with men. A related, but distinct, explanation for the different signs of the within-group peer effects for higher-education and lower-education women is that lower education

women also received less training from employers and acquired less firm-specific human capital, making them less costly to replace. If the costs to firms of keeping a position open increase for women in both subgroups when more women in the same subgroup take leave at the same time, then this difference in replacement costs across subgroups would lead to greater employment risk associated with fertility for low-education women (when more of their peers have children) and more negative peer effects. Indeed, it is well known that education lowers the risk of unemployment for women as well as men (e.g., Mincer 1991). This channel is also supported in our data. In our sample, we find that more educated women are substantially more likely to have positive earnings in 2005, and that their likelihood of having earnings in 2005 is less responsive to their own fertility, and to the interaction between their fertility and their coworkers' fertility during the sample period.²⁹

Table 7 has the same structure as Table 6, but reports estimates for peer effects by age group. As in the previous table, there is evidence of important heterogeneity, and of some positive peer effects. However, the pattern in this table differs from the pattern by education in that here the within-group effects are always negative and the cross-group effects are always positive. As in the previous table, this pattern is most consistent with the presence of positive peer effects (from social norms or coordination of childcare or information sharing) that are offset within subgroups of more similar women who are more likely to be competing (for promotions or for job security). The fact that the overall within-group peer effects are negative for both age groups, even though the within-group effect is positive for high-education women, is consistent with higher-education women being a minority within both age groups. Hence, lower-education women are likely driving the overall within-group estimates. The positive cross-group effects seem to indicate that older and younger women are not competing as strongly against one another.

In both Tables 6 and 7, the presence of heterogeneous peer effects (and of some positive peer effects) raises the potential for multiple equilibria. We compute the share of all workplaces in which the equilibrium outcome is non-uniquely determined and report these values in the tables. Multiplicity is common in our sample: over 40 percent of workplaces in Table 6 and over 50 percent in Table 8 have values of the observable controls (and unobservable errors) that would indicate more than one potential outcome. The frequency of multiple equilibria in the models with heterogeneous peer effects highlights the value of our estimation approach (that explicitly allows for these cases and uses them to infer bounds on the parameters). It also makes the discussion of marginal effects more complicated. Although we incorporated all of the potential equilibrium outcomes in estimation without selecting among them, we apply an equilibrium selection rule in computing marginal effects.

Tables 8 and 9 report predicted fertility rates and average marginal effects of covariates for women in our sample using values from the highest (Columns 1–3) or lowest (Columns 4–6)

²⁹The rates of non-participation in 2005 are 8 percent for lower-education women and 5 percent for higher-education women. Having a child during the sample period predicts a significant 4 percentage points increase in non-participation for lower-education women but an insignificant 0.2 percentage points increase for higher-education women. Finally, the estimated effect of fertility on non-participation increases for low-education women by 1.6 percentage points (from a base of 1.9 percent) for each peer who has a child, but is unaffected by peer fertility for high-education women. These patterns are consistent with the hypothesized channel, but it is important to note that the estimates do not address the endogeneity of fertility and are thus not likely to capture the true causal effects of fertility on participation.

fertility equilibrium. Following the structure of Table 5 that reports average marginal effects in the homogeneous peer effects model, the first panel shows marginal effects of covariates when peer effects have their estimated values (estimates in Table 6 are used for Table 8 and in Table 7 for Table 9) and the second panel shows marginal effects when the endogenous (strategic) peer effects are all set to zero. The maximum and minimum fertility outcomes only differ in cases of multiple equilibria. Naturally, they are identical in the bottom panel with no peer effects. As in the models with homogeneous peer effects (in Table 5), the presence of peer effects tends to reduce the impact of other factors on fertility. Going from the top to the bottom panels in Tables 8 and 9, the marginal effects of the control variables increase in magnitude when the peer effects are turned off.

The average impact of peer effects on fertility rates for women in the sample is computed by comparing the fertility rates with and without peer effects for the case of experience equals zero (when all of the women in the sample are assigned their observed values for the controls). In the model with two types of peer effects by education in Table 8, the average combination of peer effects experienced by women in the sample increases fertility rates by 25 percentage points when we resolve multiple equilibria by selecting the one with the highest fertility, but decreases rates by 7.5 percentage points if we select the lowest fertility equilibrium. In the model with two types by age in Table 9, the average effect of peer fertility has a similar impact of increasing fertility propensity by 24 percentage points when the highest fertility equilibrium is selected but of decreasing fertility by 13.9 percentage points when the lowest fertility equilibrium is selected. Although the overall peer effects were negative in the one-type (homogeneous effects) model in Table 4, the two-type estimates in Tables 6 and 7 each contained both positive and negative peer effects. The results in Table 8 and 9 show that average fertility impact of turning on these conflicting effects can be positive or negative depending on the equilibrium selection rule (choosing the maximum or minimum fertility equilibrium). This result highlights the importance of allowing for multiple equilibria in estimation without imposing arbitrary selection rules.

As discussed in Section II.A, the true equilibrium selection function is unknown and could be a function of observed and unobserved variables. Nevertheless, we can compare the distributions of women having children across workplaces predicted under each of the two extreme rules to the actual distribution in the data to see if the data support one or the other of the extreme selection rule. The predicted distributions of births under the maximum fertility (panel A) and minimum fertility (panel B) rules are shown in Appendix Figures A1 (for the heterogeneous effects by education) and A2 (by age). Compared to the actual distribution in Figure 2, it is clear that the minimum fertility rule under-predicts births (with nearly all of the mass below two births per workplace), while the maximum fertility rule over-predicts births (with sizable mass at ten or more births). The fact that actual fertility is between the two extremes is consistent with workplaces varying in their rules for selecting an equilibrium when faced with two or more potential equilibria. It is also consistent with workplaces sometimes selecting an equilibrium with an intermediate fertility level when three or more equilibria are available.

5. Policy Simulations: Sorting of Women across Establishments

This section reports results from simulation exercises that explore the importance of sorting of women across establishments on overall fertility rates. As discussed above, one implication of peer effects in fertility among coworkers is that the female labor force participation would then affect fertility, not just through the direct effects of working for pay on the opportunity costs of childbearing, but also through interactions at the workplace. In that case, *where* each woman works, and the peers *with whom* she works, can also affect fertility decisions.

In our policy experiment, we measure the implied effect of sorting on fertility rates using parameter estimates from our structural models. We do this by creating a set of hypothetical, equally-sized work establishments, and assigning 7 women from our estimation sample to these establishments (thereby creating peer groups) according to two extremes of sorting. First, in what we call “perfect” sorting, we order all women according to the individual propensity to have a child in the sample period, based only on exogenous individual characteristics and excluding the endogenous and contextual peer effects, and then assign them to workplaces in order (the first 7 women are assigned to the first workplace, the next 7 to the next, etc.). Our second assignment rule, called “random”, involves filling each open slot in order with a woman drawn randomly (without replacement) from the estimation sample. To explore the interaction between establishment size and sorting, we also create additional samples of simulated workplaces with 15 women at each workplace.³⁰

Table 10 summarizes the different predicted fertility rates for each of these samples using estimates from structural models with homogeneous and heterogeneous peer effects. In cases of multiple equilibria, we report equilibrium fertility rates using either the highest or the lowest equilibrium fertility rate. The first row reports the predicted fertility rates for each of the simulated sample from a model with endogenous peer effects set to zero. This is used as a benchmark for computing incremental endogenous peer effects. These predictions incorporate contextual peer effects, which is why the sorting rule affects fertility even in the baseline.

Fertility rates decline in all four samples when the homogeneous peer effects are applied in the next row. The magnitude of the fertility decline (both in absolute terms and as a percent of the baseline rate) is larger for the case of perfect sorting. Under perfect sorting, women who are unlikely to have children (absent peer effects) are those who experience the *smallest* negative peer effects while those who are most likely to have children experience the *largest* peer effects because more of their coworkers have children. In the case of negative peer effects, the concentration of women with higher fertility preferences together leads to lower fertility rates than the rates that would result if those women were most dispersed across workplaces.

³⁰We develop our policy simulations to focus on composition of the workplace instead of variation in its size. We use 7 because it is the median size and 15 to capture the 90th percentile. Appendix Table A4 shows that the average levels of integration of women by age and experience across establishments in our sample lie between the levels of integration for the random and sorted samples.

The impact of sorting is more complex in the case of heterogeneous peer effects. The results for the two-type case by age group are reported next in Table 10. As described in the previous section (and shown in Table 7), the within-type effects are negative for both types but the cross-type effects are both positive. This means that different sorting rules not only affect the likelihood that each individual woman has a peer with a strong taste for fertility but that the sorting rules also affect the chance that she has a peer of her same type (or opposite type) with such tastes. Under perfect sorting, women have mostly same-type peers, and experience mostly negative within-type peer effects. With random sorting, women face a mix of peers and can experience both positive and negative peer effects. The result is that sorting has a much larger impact on fertility in the case with heterogeneous peer effects and also a much larger effect on the incremental effect of workplace peer interactions.

When we select the maximum fertility equilibrium, we find that sorting does more than simply moderate the size of the peer effect; it reverses the direction of the overall peer effect, going from large and negative with perfect sorting (where the negative within-type effects dominate) to large and positive with random sorting (where more positive cross-type effects are present). Appendix Table A3 shows that the 91 percent decline in fertility attributable to peer effects under perfect sorting by fertility propensity is repeated if we instead apply a scheme that first sorts women by age and then by fertility propensity or a scheme that only sorts based on age. This supports the interpretation that workplace segregation by age, and greater exposure to peers with negative spillovers, is the main reasons that sorting by fertility propensity leads to a reversal in the sign of the overall effect of peer interactions on fertility rates.

This pattern is identical between workplaces with 7 and 15 women, but the positive overall peer effect with random sorting is no longer present if we select the minimum fertility outcome. Instead, shifting from perfect to random sorting leaves the negative peer effects relatively unchanged. The choice of equilibrium selection rule has an even larger impact on the average peer effect in the model with heterogeneous effects by education. When we select the equilibrium with the highest fertility rate, peer effects lead to higher fertility rates (as in Table 8), and the positive effect is substantially larger with random peers (161 percent) compared to perfect sorting (91 percent) or sorting by type, with or without also sorting by fertility propensity (also 91 percent; Appendix Table A3). When we select the lowest fertility equilibrium, the peer effects lead to the same 64 percent decline in fertility with random or matched peers. Overall, the implication is that greater workplace segregation by age or education leads to lower overall fertility.

The fact that predicted fertility rates are consistently higher when women are randomly distributed across establishments suggests a secondary mechanism through which increasing the generosity of mandated family leave (or other family friendly policies) could raise fertility rates for working women. In addition to the direct effect on individual women, such policies could *reduce* the degree of sorting of women according to fertility tastes, because women with high fertility tastes will not cluster as much at relatively family friendly workplaces, which would tend to raise fertility rates. This is especially true for women in the private sector, where family-friendly policies are less generous and more heterogeneous and where workplaces are more segregated by age and education (see Table A4). It is worth

noting that this change may even reflect an efficiency improvement if the more even distribution of women having children across workplaces lowers the total cost to employers and co-workers from fertility spillovers (possibly caused by scale diseconomies in leave taking) implied by the negative peer effects.

6. Conclusion

Using register-based data on the population of Denmark, and a sample of women working at medium-sized establishments in November 2002, this paper finds evidence that interactions between female coworkers generate substantial peer effects in fertility outcomes over 2003–2005. The peer effects are negative overall, but they are also heterogeneous across worker types defined by age or education. Cross-type effects are positive in all cases, suggesting either positive social effects or scale economies in childcare. Within type effects are negative for low-education women (about two-thirds of the sample) and for both older and younger women. The negative effects are likely due to career concerns related to having a child when several coworkers are doing the same – if employers face sufficiently large costs, they may not be able to retain all of the workers through their maternity leaves or to provide all returning mothers with the same opportunities for advancement. The pattern of these relationships suggests that economic interactions among coworkers affect individual fertility decisions. The policy simulation exercises show that the peer effects (both endogenous and contextual) imply that the distribution of where women work can affect overall fertility rates.

This paper is the first to study fertility decisions between coworkers using a game theory model for strategic interactions. Our empirical approach may be useful in studying peer interactions in other contexts. A limitation of the current approach is that the net utility function upon which the agents in our model base their fertility decisions is itself a reduced form representation of a more complex process. For example, women make choices concerning both their fertility and labor supply, and each of these decisions may be affected by the woman's age and education. This means that the estimated effects of these variables in the current paper should be not interpreted as only capturing shifters of fertility preferences, but also shifts in anticipated income effects from fertility. A natural extension for future work would be to model and estimate the conjoint strategic decisions of fertility and subsequent labor supply among coworkers.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

References

- Adda, J.; Dustmann, C.; Stevens, K. IZA Discussion Paper No. 6201. 2011. The Career Costs of Children.
- Bandiera O, Barankay I, Rasul I. Social Incentives in the Workplace. *Review of Economic Studies*. 2010; 77:417–458.
- Becker, G. An Economic Analysis of Fertility. In: Becker, G., editor. *Demographic and Economic Change in Developed Countries*. Princeton, N.J: Princeton University Press; 1960.
- Bernheim BD. A Theory of Conformity. *Journal of Political Economy*. 1994; 102:841–877.

- Berry S, Levinsohn J, Pakes A. Automobile Prices in Market Equilibrium. *Econometrica*. 1995; 63:841–890.
- Berry, S.; Reiss, P. Empirical Models of Entry and Market Structure. In: Armstrong, M.; Porter, R., editors. *Handbook of Industrial Organization*. Vol. 3. Elsevier; 2007. p. 1845-1886.
- Berry, S.; Tamer, E. Blundell, Newey and Persson. Identification in Models of Oligopoly Entry. *Advances in Economics and Econometrics; Ninth World Congress; Cambridge University Press*; 2006. p. 46-85.
- Blake D, Mayhew L. On The Sustainability of the UK State Pension System in the Light of Population Ageing and Declining Fertility. *Economic Journal*. 2006; 116:F286–F305.
- Bloom, D.; Canning, D.; Gunther, I.; Linnemayr, S. Harvard School of Public Health Working Paper. 2008. Social Interactions and Fertility in Developing Countries.
- Blume, L.; Durlauf, S. Identifying Social Interactions: A Review. San Francisco: Jossey-Bass; 2005.
- Borsch-Supan A. A Model Under Siege: A Case Study of the German Retirement Insurance System. *Economic Journal*. 2000; 110:F24–F45.
- Bresnahan TF, Reiss PC. Entry in Monopoly Market. *Review of Economic Studies*. 1990; 57:531–553.
- Brock, WA.; Durlauf, SN. Interactions-based Models. In: Heckman; Leamer, editors. *Handbook of Econometrics*. Vol. 5. Elsevier; 2001. p. 3297-3380.
- Brock WA, Durlauf SN. Identification of Binary Choice Models with Social Interactions. *Journal of Econometrics*. 2007; 140:52–75.
- Card D, Giuliano L. Peer Effects and Multiple Equilibria in the Risky Behavior of Friends. *Review of Economics and Statistics*. 2013; 95:1130–1149.
- Chernozhukov V, Hong H, Tamer E. Estimation and Confidence Regions for Parameter Sets in Econometric Models. *Econometrica*. 2007; 75:1243–1284.
- Ciliberto F, Tamer E. Market Structure and Multiple Equilibria in Airline Markets. *Econometrica*. 2009; 77:1791–1828.
- Cohen A, Freeborn B, McManus B. Competition and Crowding-Out in the Market for Outpatient Substance-Abuse Treatment. *International Economic Review*. 2013; 54:159–184.
- Cotter, DA.; Hermesen, JM.; Vanneman, R. Gender inequality at work. Russell Sage Foundation; 2004.
- Crane J. The epidemic theory of ghettos and neighborhood effects on dropping out and teenage childbearing. *American Journal of Sociology*. 1991; 96:1226–1259.
- Dahl GB, Løken KV, Mogstad M. Peer Effects in Program Participation. *American Economic Review*. 2014; 104:2049–74.
- Duflo E, Saez E. The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence From a Randomized Experiment. *Quarterly Journal of Economics*. 2003; 118:815–842.
- Galor, Oded. The Demographic Transition and the Emergence of Sustained Economic Growth. *Journal of the European Economic Association*. 2005; 3:494–504.
- Graham BS, Hahn J. Identification and Estimation of the Linear-in-Means Model of Social Interactions. *Economics Letters*. 2005; 88:1–6.
- Grant, J.; Hoorens, S.; Sivadasan, S.; van het Loo, M.; DaVanzo, J.; Hale, L.; Gibson, S.; Butz, W. RAND Corporation Technical report. 2004. Low Fertility and Population Aging: Causes, Consequences and Policy Options.
- Grieco P. Discrete games with flexible information structures: an application to local grocery markets. *RAND Journal of Economics*. 2014; 45:303–340.
- Hensvik, L.; Nilsson, P. Working Paper. Institute for Evaluation of Labour Market and Education Policy; 2010. Businesses, buddies and babies: social ties and fertility at work.
- Hotz, JV.; Klerman, J.; Willis, R. The economics of fertility in developed countries. In: Rosenzweig; Stark, editors. *Handbook of Population and Family Economics*. Vol. 1. Elsevier; 1997. p. 275-347.
- Jones, L.; Schoonbroodt, A.; Tertilt, M. Fertility Theories: Can They Explain the Negative Fertility-Income Relationship?. In: Shoven, editor. *Demography and the Economy*. University of Chicago Press; 2011.
- Juul, Karmaus SW, Olsen J. The European Infertility and Sub-fecundity Study Group. Regional Differences in Waiting Time to Pregnancy: Pregnancy-based Surveys from Denmark, France, Germany, Italy and Sweden. *Human Reproduction*. 1999; 14:1250–1254. [PubMed: 10325272]

- Kearney MS. Is There an Effect of Incremental Welfare Benefits on Fertility Behavior? A Look at the Family Cap. *Journal of Human Resources*. 2004; 39:295–325.
- Kohler, H-P. *Fertility and Social Interaction - An Economic Perspective*. Oxford: Oxford University Press; 2001.
- Krauth BV. Simulation-based Estimation of Peer Effects. *Journal of Econometrics*. 2006; 133:243–271.
- Kravdal O. Education and Fertility in Sub-saharan Africa: Individual and Community Effects. *Demography*. 2002; 39:233–250. [PubMed: 12048950]
- Kuziemko, I. Mimeo. Harvard University; 2006. Is Having Babies Contagious? Estimating Fertility Peer Effects Between Siblings.
- Lazear EP, Rosen S. Rank-order Tournaments as Optimum Labor Contracts. *Journal of Political Economy*. 1981; 89:841–864.
- Lindh T, Malmberg B. Age structure effects and growth in the OECD 1950–1990. *Journal of Population Economics*. 1999; 12:431–449.
- Lovenheim MF, Mumford K. Do Family Wealth Shocks Affect Fertility Choices? Evidence from the Housing Market. *Review of Economics and Statistics*. 2013; 95:464–475.
- Manski C. Identification of Endogenous Social Effects: The Reflection Problem. *Review of Economic Studies*. 1993; 60:531–542.
- Manski, C. *Identification Problems in the Social Sciences*. Harvard University Press; 1995.
- Manski C, Mayshar J. Private incentives and social interactions: fertility puzzles in Israel. *Journal of the European Economic Association*. 2003; 1:181–211.
- Mas A, Moretti E. Peers at Work. *American Economic Review*. 2009; 99:112–145.
- Miller A. The Effects of Motherhood Timing on Career Path. *Journal of Population Economics*. 2011; 24:1071–1100.
- Mincer, J. NBER Working Paper 3837. 1991. Education and Unemployment of Women.
- Moffitt, R. The Effect of Welfare on Marriage and Fertility: What Do We Know and What Do We Need to Know?. National Research Council; 1998.
- Montgomery, MR.; Casterline, JB. Population Council Policy Research Division Working Paper No. 119. 1998. Social Networks and the Diffusion of Fertility Control.
- Montgomery MR, Casterline JB. Social Learning, Social Influence, and New Models of Fertility. *Population and Development Review*. 1996; 22:151–175. (Supplement: Fertility in the United States: New Patterns, New Theories).
- Mumford, K. SIEPR Discussion Paper No. 0620. 2007. The Optimal Tax Treatment of Families with Children.
- Murray C. Welfare and the Family: The U.S. Experience. *Journal of Labor Economics*. 1993; 11:S244–S262.
- Nielsen HS, Simonsen M, Verner M. Does the Gap in Family-friendly Policies Drive the Family Gap? *Scandinavian Journal of Economics*. 2004; 106:721–744.
- Ponomareva M, Tamer E. Misspecification in moment inequality models: back to moment equalities? *Econometrics Journal*. 2011; 14:186–203.
- Rangel A. Forward and Backward Intergenerational Goods: Why Is Social Security Good for the Environment? *American Economic Review*. 2003; 93:813–834.
- Rennhoff AD, Owens MF. Competition and the Strategic Choices of Churches. *American Economic Journal: Microeconomics*. 2012; 4:152–170.
- Rosse J. Estimating Cost Function Parameters without Using Cost Data: Illustrated Methodology. *Econometrica*. 1970; 38:256–75.
- Sacerdote B. Peer Effects with Random Assignment: Results for Dartmouth Roommates. *Quarterly Journal of Economics*. 2001; 116:681–704.
- Schmalensee, R. Inter-industry studies of structure and performance. In: Schmalensee, Richard; Willig, Robert, editors. *Handbook of Industrial Organization*. Vol. 2. Elsevier Press; 1989. p. 951-1009. Chapter 16
- South S, Baumer E. Deciphering community and race effects on adolescent premarital childbearing. *Social Forces*. 2000; 78:1379–1408.

- Tamer E. Incomplete Simultaneous Discrete Response Model with Multiple Equilibria. *Review of Economic Studies*. 2003; 70:147–165.
- Train, K. *Discrete Choice Methods with Simulation*. Cambridge University Press; 2003.
- Waldfogel J. Understanding the ‘Family Gap’ in Pay for Women with Children. *Journal of Economic Perspectives*. 1998; 12:137–156.
- Willis R. A New Approach to the Economic Theory of Fertility Behavior. *Journal of Political Economy*. 1973; 81(S14–S64)

Figure 1.
Distribution of Women (Ages 20–40) across Workplaces, November 2002

Figure 2.
Distribution of Women with Births across Workplaces, 2003–2005

Figure 3.
Distribution of Share of Women with Births across Workplaces, 2003–2005

Table 1

Illustration of Nash Equilibrium Outcomes with Two Agents

Index of possible outcomes	Women having children	v_{1e}	v_{2e}
1	0	$\alpha_r X_{1e} + \varepsilon_{1e}^{(s)} < 0$	$\alpha_r X_{1e} + \varepsilon_{1e}^{(s)} < 0$
2	1	$\alpha_r X_{1e} + \varepsilon_{1e}^{(s)} > 0$	$\alpha_r X_{2e} + \varepsilon_{2e}^{(s)} + \delta_r^r < 0$
2	1	$\alpha_r X_{1e} + \varepsilon_{1e}^{(s)} + \delta_r^r < 0$	$\alpha_r X_{2e} + \varepsilon_{2e}^{(s)} > 0$
3	2	$\alpha_r X_{1e} + \delta_r^r + \varepsilon_{1e}^{(s)} > 0$	$\alpha_r X_{2e} + \delta_r^r + \varepsilon_{2e}^{(s)} > 0$

Notes: This table illustrates the case of a workplace with two agents and three possible equilibria. Each row shows the conditions on the net utility of each agent that would make that outcome a valid Nash Equilibrium. There are two sets of conditions that would each make one woman having a child (outcome #2) an equilibrium outcome.

Table 2

Descriptive Statistics

<i>Panel A -- Individuals</i>			
	No Birth	Birth	Total
Birth 2003–2005	0	1	0.237
Age > 30	0.464	0.595	0.495
College educated	0.269	0.401	0.300
Experience (in decades)	0.599	0.589	0.597
No children before 2003	0.507	0.514	0.509
One child (born before 2003)	0.124	0.344	0.176
Two children (born before 2003)	0.282	0.123	0.244
More than two children (born before 2003)	0.087	0.018	0.071
Sister fertility (any sister with birth in 2000–2002)	0.118	0.141	0.123
Number of observations	115,541	35,953	151,494
<i>Panel B -- Establishments</i>			
	Mean	Std. Dev.	
Establishment size	20.46	10.11	
- Share 20 employees or fewer	0.63		
- Share 21–30 employees	0.20		
- Share 31–40 employees	0.10		
- Share more than 40 employees	0.07		
Number of women aged 20–40	4.78	3.83	
One woman aged 20–40	0.18		
Number of women aged 20–40 who have a child in 2003–2005	1.13	1.47	
Private firm	0.72		
# Establishments	31,725		

Source: Administrative register-based data on the population of Denmark. Sample of establishments is limited to those with 10 to 50 employees. Sample of individuals is limited to women aged 20 to 40 who worked at one of these establishments in November 2002. Peer groups are defined based on work establishment in November 2002.

Table 3

Probit Estimates Relating Own Fertility to Peer Fertility, 2003–2005

	# Peers Having Children	
	Marginal Effect	S.E.
Exploratory Conditioning Sets		
None	0.0103***	0.0071
+ Individual characteristics	0.0033***	0.0007
+ Sister fertility	0.0032***	0.0007
+ Relationship status and partner characteristics	0.0032***	0.0007
+ Establishment characteristics	0.0030***	0.0008
+ Peer characteristics	0.0037***	0.0008
Main Conditioning Sets		
Basic	0.0051***	0.0007
+ Contextual effects	0.0055***	0.0007
Main Conditioning Sets with Random Effects		
Basic	−0.0943***	0.0011
+ Contextual effects	−0.1013***	0.0012

Notes: Robust standard errors clustered at the work establishment level. The first set of estimates is from separate Probit models that add an extensive set of covariates. The full list of exploratory conditioning variables is available from the authors. The second set of estimates uses the conditioning variable sets in the main models in Tables 4 to 10. The third set uses these same conditioning variables but also includes establishment-specific random effects (drawn from a normal distribution with a mean of zero and variance of 1). The total number of observations is 151,494.

Table 4**Fertility Game Estimates with Homogeneous Peer Effects**

Variable	(1)	(2)	(3)	(4)
Peer fertility	[0.044, 0.064]	[−0.307, −0.252]	[−0.278, −0.272]	[−0.275, −0.271]
<i>Individual controls:</i>				
Constant	[−1.701, −1.670]	[−2.204, −2.064]	[−2.405, −2.253]	[−2.164, −2.157]
Age > 30	[0.612, 0.702]	[1.014, 1.203]	[1.176, 1.332]	[1.303, 1.309]
College educated	[0.419, 0.458]	[0.889, 1.040]	[0.989, 1.042]	[0.878, 0.884]
Experience	[0.466, 0.582]	[0.852, 0.095]	[0.964, 0.994]	[0.902, 0.906]
One child	[0.728, 0.814]	[0.943, 1.130]	[0.979, 1.072]	[1.252, 1.261]
Two children	[−0.215, −0.137]	[−0.830, −0.657]	[−0.907, −0.706]	[−1.035, −1.024]
<i>Instrument:</i>				
Sister fertility			[0.190, 0.255]	[0.132, 0.145]
<i>Contextual effects:</i>				
Median age >30				[−0.231, −0.223]
Median educated				[0.011, 0.019]
Median experience				[0.034, 0.044]
Median one child				[−0.611, −0.599]
Median two children				[0.1010.108]
Firm and Establishment Random Errors	N	Y	Y	Y
Function value	1.57E+04	1.60E+04	1.65E+04	1.64E+04
Fit	0.1277	0.1245	0.1257	0.1265
Percent multiple	0	0	0	0

Notes: Each column reports 95-percent confidence bounds on parameter values for the effects of variables on the fertility decision in different models of the fertility game. Column (1) is from the basic model with only individual controls; column (2) adds firm-specific and establishment-specific random errors; column (3) adds the control for sister fertility; column (4) adds contextual effects summarizing coworker characteristics. The total number of observations is 31,725.

Table 5**Marginal Effects in the Homogeneous Peer Effects Model**

Variable	1	0	Marginal Effect
<i>Peer Effects On</i>			
Age > 30	0.2966	0.2112	0.0854
Education	0.2832	0.2264	0.0568
One child	0.3095	0.2197	0.0898
Two children	0.204	0.2671	-0.0631
Sister fertility	0.2456	0.2371	0.0085
Experience	0.2435	0.2381	0.0054
<i>Peer Effects Off</i>			
Age > 30	0.3542	0.2583	0.0959
Education	0.3369	0.2738	0.0631
One child	0.3679	0.2678	0.1001
Two children	0.2512	0.323	-0.0718
Sister fertility	0.2959	0.2863	0.0096
Experience	0.2936	0.2875	0.0061
Marginal effect of peer fertility			-0.0494

Notes: The values in this table are predicted population average fertility rates for different values of the explanatory variables in models with and without peer effects in fertility. The first column has fertility rates with the binary variables (for age, education, one child, two children, and sister fertility) set to one for each woman and the second column has the same variables set to zero. For the discrete variable measuring experience, the second column has the actual value and the first column increases the variable by one unit (corresponding to 10 years). The marginal effect in the third column is the difference in fertility with the variable set to one versus zero. The average impact of peer effects on fertility in the sample is shown by the difference in predicted fertility with peer effects on versus off in the case of experience.

Table 6**Fertility Game Estimates by Education**

Variable	(1)	(2)	(3)	(4)
<i>Peer Effects by Education Level:</i>				
Low on low	[-0.791, -0.753]	[-0.803, -0.741]	[-0.823, -0.722]	[-1.140, -1.062]
Low on high	[1.7901, 1.812]	[1.838, 1.907]	[1.660, 1.741]	[1.953, 2.018]
High on low	[6.337, 6.651]	[5.998, 6.088]	[6.282, 10.574]	[10.460, 10.525]
High on high	[5.994, 6.619]	[6.602, 6.693]	[8.645, 9.906]	[9.176, 9.274]
<i>Individual controls:</i>				
Constant	[-2.512, -2.300]	[-2.824, -2.762]	[-2.719, -2.471]	[-2.500, -2.169]
Age > 30	[1.018, 1.086]	[1.219, 1.290]	[0.700, 0.746]	[0.852, 1.001]
Education	[-1.907, -1.744]	[-1.841, -1.748]	[-0.157, 0.151]	[-0.183, -0.105]
Experience	[1.005, 1.345]	[1.394, 1.476]	[0.301, 1.403]	[0.315, 0.482]
One Child	[0.796, 0.891]	[0.889, 0.952]	[1.259, 1.600]	[1.139, 1.322]
Two Children	[-0.327, -0.229]	[-0.157, -0.095]	[0.467, 1.071]	[-0.359, 0.180]
<i>Instrument:</i>				
Sister fertility		[0.068, 0.156]	[0.085, 0.202]	[0.083, 0.234]
<i>Contextual effects:</i>				
Median age >30			[0.367, 0.513]	[0.131, 0.512]
Median educated			[-2.196, -1.470]	[-3.971, -3.856]
Median experience			[-0.022, 0.615]	[0.554, 0.852]
Median one child			[-0.435, -0.302]	[-0.593, -0.293]
Median two children			[-0.967, -0.546]	[-0.495, -0.182]
Correlation				[0.372, 0.495]
Function value	1.47E+04	1.52E+04	1.50E+04	1.50E+04
Fit	0.1182	0.1188	0.1186	0.1194
Share multiple	0.4234	0.4264	0.4295	0.4384

Notes: The 95-percent confidence parameter bounds in this table are from expanded versions of the models in Table 4 that allow for different effects of peer fertility that depend on the education level of the agent and the coworker. All models include firm-specific and establishment-specific error terms. Column (1) has individual controls; column (2) adds the instrument based on sisters' fertility; column (3) adds contextual effects; column (4) adds correlated error terms within work establishments. The total number of observations is 31,725.

Table 7

Fertility Game Estimates by Age

Variable	(1)	(2)	(3)	(4)
<i>Peer Effects by Age:</i>				
Younger on younger	[-2.834, -2.768]	[-3.295, -3.248]	[-2.163, -2.122]	[-3.760, -3.702]
Younger on older	[3.846, 3.906]	[4.296, 4.339]	[3.361, 3.420]	[4.469, 4.542]
Older on younger	[4.088, 4.120]	[5.436, 5.507]	[5.638, 5.679]	[5.510, 5.575]
Older on older	[-3.962, -3.853]	[-5.272, -5.210]	[-5.698, -5.661]	[-6.417, -6.344]
<i>Individual controls:</i>				
Constant	[-2.013, -1.956]	[-2.110, -2.063]	[-2.265, -2.231]	[-2.100, -2.040]
Age > 30	[0.495, 0.558]	[0.570, 0.620]	[-0.107, -0.075]	[-0.123, -0.067]
Education	[0.392, 0.463]	[0.327, 0.376]	[0.288, 0.320]	[0.599, 0.645]
Experience	[0.523, 0.594]	[0.555, 0.599]	[0.077, 0.118]	[-0.169, -0.114]
One Child	[0.674, 0.729]	[0.698, 0.749]	[0.921, 0.961]	[1.064, 1.122]
Two Children	[-0.449, -0.384]	[-0.370, -0.305]	[-0.016, 0.045]	[0.405, 0.475]
<i>Instrument:</i>				
Sister fertility		[0.135, 0.183]	[0.098, 0.132]	[0.103, 0.202]
<i>Contextual effects:</i>				
Median age >30			[0.955, 0.997]	[0.991, 1.050]
Median educated			[0.216, 0.253]	[-0.119, -0.073]
Median experience			[0.534, 0.602]	[0.701, 0.779]
Median one child			[-0.122, -0.087]	[-0.400, -0.341]
Median two children			[-0.306, -0.244]	[-0.662, -0.602]
Correlation				[0.701, 0.763]
Function value	1.37E+04	1.41E+04	1.39E+04	1.39E+04
Fit	0.117	0.1171	0.1178	0.1167
Share multiple	0.5440	0.5540	0.5504	0.5665

Notes: The 95-percent confidence parameter bounds in this table are from expanded versions of the models in Table 4 that allow for different effects of peer fertility that depend on the age group of the agent (over or under 30 years of age) and the coworker. All models include firm-specific and establishment-specific error terms. Column (1) has individual controls; column (2) adds the instrument based on sisters' fertility; column (3) adds contextual effects; column (4) adds correlated error terms within work establishments. The total number of observations is 31,725.

Table 8

Marginal Effects in the Peer Effects Model with Two Types by Education

Variable	Maximum Fertility Equilibrium		Minimum Fertility Equilibrium	
	1	0	1	0
<i>Peer Effects On</i>				
Age > 30	0.4801	0.4309	0.0492	0.1028
Education	0.4490	0.4487	0.0003	0.1194
Experience	0.4524	0.4508	0.0016	0.1216
One child	0.4930	0.4407	0.0523	0.1117
Two children	0.4610	0.4470	0.0140	0.1181
Sister fertility	0.4594	0.4485	0.0109	0.1194
<i>Peer Effects Off</i>				
Age > 30	0.2325	0.1700	0.0625	0.1700
Education	0.1963	0.1917	0.0046	0.1917
Experience	0.1979	0.1961	0.0018	0.1961
One child	0.2453	0.1838	0.0615	0.1838
Two children	0.2087	0.1907	0.0180	0.1907
Sister fertility	0.2091	0.1920	0.0171	0.1920
Marginal effect of peer fertility			0.2547	-0.0745

Notes: The values in this table are predicted population average fertility rates for different values of the explanatory variables in the models with peer effects by education group and without peer effects in fertility. In cases of multiple equilibria, the maximum fertility equilibrium is the equilibrium with the largest number of births and the minimum fertility equilibrium is the one with the smallest number (these are the same with no peer effects). The first column has fertility rates with the binary variables (for age, education, one child, two children, and sister fertility) set to one for each woman and the second column has the same variables set to zero. For the discrete variable measuring experience, the second column has the actual value and the first column increases the variable by one unit (corresponding to 10 years). The marginal effect in columns 3 and 4 are differences in fertility with the variable set to one versus zero. The average impacts of peer effects on fertility in the sample are shown by the difference in predicted fertility with peer effects on versus off in the case of experience.

Table 9

Marginal Effects in the Peer Effects Model with Two Types by Age

Variable	Maximum Fertility Equilibrium		Minimum Fertility Equilibrium		Marginal Effect
	1	0	1	0	
Peer Effects On					
Age > 30	0.5029	0.5068	0.1207	0.1247	-0.0040
Education	0.5293	0.4996	0.1489	0.1169	0.0320
Experience	0.5047	0.5050	0.1225	0.1229	-0.0004
One child	0.5503	0.4952	0.1694	0.1131	0.0563
Two children	0.5205	0.4996	0.1427	0.1172	0.0255
Sister fertility	0.5131	0.5041	0.1318	0.1219	0.0099
Peer Effects Off					
Age > 30	0.2579	0.2647	0.2579	0.2647	-0.0068
Education	0.2975	0.2524	0.2975	0.2524	0.0451
Experience	0.2619	0.2618	0.2619	0.2618	0.0001
One child	0.3266	0.2479	0.3266	0.2479	0.0787
Two children	0.2889	0.2510	0.2889	0.2510	0.0379
Sister fertility	0.2802	0.2587	0.2802	0.2587	0.0215
Marginal effect of peer fertility					-0.1389

Notes: The values in this table are predicted population average fertility rates for different values of the explanatory variables in models with by peer effects by age group and without peer effects in fertility. In cases of multiple equilibria, the maximum fertility equilibrium is the equilibrium with the largest number of births and the minimum fertility equilibrium is the one with the smallest number (these are the same with no peer effects). The first column has fertility rates with the binary variables (for age, education, one child, two children, and sister fertility) set to one for each woman and the second column has the same variables set to zero. For the discrete variable measuring experience, the second column has the actual value and the first column increases the variable by one unit (corresponding to 10 years). The marginal effect in columns 3 and 6 are differences in fertility with the variable set to one versus zero. The average impacts of peer effects on fertility in the sample are shown by the difference in predicted fertility with peer effects on versus off in the case of experience.

Table 10

Impact of Alternative Rules for Assigning Women to Establishments

Establishment size	7 Women		15 Women	
Sorting and Assignment Rule	Perfect	Random	Perfect	Random
Homogeneous peer effects	0.192	0.228	0.191	0.227
Baseline: No endogenous peer effects	0.281	0.308	0.281	0.308
Incremental peer effect	−0.089	−0.080	−0.090	−0.081
Percent of baseline	−32%	−26%	−32%	−26%
Heterogeneous peer effects by age				
Baseline: No endogenous peer effects	0.256	0.302	0.256	0.301
Maximum fertility equilibrium	0.064	0.738	0.064	0.738
Incremental peer effect	−0.192	0.436	−0.192	0.437
Percent of baseline	−75%	145%	−75%	145%
Minimum fertility equilibrium	0.064	0.075	0.061	0.075
Incremental peer effect	−0.192	−0.227	−0.195	−0.226
Percent of baseline	−75%	−75%	−76%	−75%
Heterogeneous peer effects by education				
Baseline: No endogenous peer effects	0.196	0.260	0.196	0.260
Maximum fertility equilibrium	0.375	0.678	0.375	0.679
Incremental peer effect	0.179	0.418	0.179	0.419
Percent of baseline	91%	161%	91%	161%
Minimum fertility equilibrium	0.070	0.094	0.070	0.094
Incremental peer effect	−0.126	−0.166	−0.126	−0.166
Percent of baseline	−64%	−64%	−64%	−64%

Notes: The baseline predicted fertility rates (shares of women having at least one child during the sample period) are from models with endogenous peer effects set to zero. The respective baselines match the distributions of individual shocks used in the one-type (uncorrelated) and two-type (correlated) models. Columns 1 and 2 are from simulated establishments with 7 fertile-age women and Columns 3 and 4 have 15 women at each establishment. The perfect sorting rule assigns women to peers that are most similar to them (in terms of their predicted fertility propensity, based on individual covariates). The random sorting rule draws women at random from the sample (without replacement) and assigns them to workplaces. In the heterogeneous peer effects models, we account for multiple equilibria by either selecting the equilibrium with the largest (maximum fertility equilibrium) or smallest (minimum) number of births. The second row reports predicted fertility rates in the model with homogeneous peer effects. The next two rows report the incremental impacts of peer effects on fertility in the homogeneous effect model, first as a change in the share of women having children (relative to the baseline) and next as a proportion of the baseline rate. The remaining rows report predicted fertility rates and incremental peer effects from models that incorporate heterogeneous effects by age or education.