

DISCUSSION PAPER SERIES

IZA DP No. 12904

**Child Care Markets, Parental Labor
Supply, and Child Development**

Samuel Berlinski
Maria Marta Ferreyra
Luca Flabbi
Juan David Martin

JANUARY 2020

DISCUSSION PAPER SERIES

IZA DP No. 12904

Child Care Markets, Parental Labor Supply, and Child Development

Samuel Berlinski

Inter-American Development Bank and IZA

Maria Marta Ferreyra

The World Bank

Luca Flabbi

University of North Carolina at Chapel Hill and IZA

Juan David Martin

The World Bank

JANUARY 2020

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

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

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

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Child Care Markets, Parental Labor Supply, and Child Development*

We develop and estimate a model of child care markets that endogenizes both demand and supply. On the demand side, families with a child make consumption, labor supply, and child-care decisions within a static, unitary household model. On the supply side, child care providers make entry, price, and quality decisions under monopolistic competition. Child development is a function of the time spent with each parent and at the child care center; these inputs vary in their impact. We estimate the structural parameters of the model using the 2003 Early Childhood Longitudinal Study, which contains information on parental employment and wages, child care choices, child development, and center quality. We use our estimates to evaluate the impact of several policies, including vouchers, cash transfers, quality regulations, and public provision. Among these, a combination of quality regulation and vouchers for working families leads to the greatest gains in average child development and to a large expansion in child care use and female labor supply, all at a relatively low fiscal cost.

JEL Classification: J13, J22, L1

Keywords: child care markets, child care quality, early childhood development, female labor supply

Corresponding author:

Luca Flabbi
Economics Department
University of North Carolina at Chapel Hill
107 Gardner Hall, CB 3305
Chapel Hill, NC 27599-3305
USA
E-mail: luca.flabbi@unc.edu

* We thank Raquel Bernal, Matias Busso, Daniela Del Boca and Chris Flinn for their useful comments. Juliana Chen Peraza provided excellent research assistance. The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the Inter-American Development Bank, The World Bank Group, their Board of Directors, or the countries they represent.

1 Introduction

In the development of a child, parental inputs play a critical role.¹ They include goods and services that can be bought in the market as well as parents' time. Both input types are intimately related to the labor market, since supplying labor in the market generates income but also lowers the time available for children. A specific institution is at the center of this trade-off in the child's early years: the child care center. Center-based care can be bought in the market to relax the constraint on parents' time. Not surprisingly, increasing access to center-based care for families with young children is key to the debate on family-friendly policies.² The objectives of those who support it are not always the same. The traditional argument is that increasing access to center-based care would raise female labor force participation, which would increase national output, favor the accumulation of human capital by women, and improve women's economic self-sufficiency.³ More recent views have focused on the impact on center-based care on children, as high-quality child care for disadvantaged children has been shown to improve children's skills and well-being, and to have positive long-term impacts.⁴

In countries with a well-developed child care market, child care use is a market equilibrium outcome. In this market, providers choose the quality and price of their services, and parents choose the quantity and quality of child care to purchase. Therefore, when evaluating proposals to expand child care access, the policymaker must investigate their impact on child care demand and supply while taking into account their effect on child development and parental labor market outcomes.

This paper develops and estimates a model of supply and demand for child care. On the demand side, parents who care about household consumption, parental leisure, and child development, and who face a set of wage offers and child care options, decide on consumption, labor supply, and child care. Child development, which we measure by a cognitive score, is a function of the time spent with each parent and at the child care center. On the supply side, we model the market of child care providers as one of monopolistic competition with vertical and horizontal differentiation. Importantly, we endogenize provider entry through a two-stage game. In the first stage, potential entrants decide on entry and quality. In the second stage, actual entrants choose prices as they face marginal costs which are increasing in quality.⁵ We model heterogeneity on both sides of the market. Households are heterogeneous in wage offers, education, initial child development, and preferences; and child care providers are heterogeneous in their marginal costs to provide each quality level (representing, for

¹The relationship between parental inputs and child development has been heavily studied in the literature. See, for instance, the reviews in Heckman et al. (2005), Björklund and Salvanes (2011), and Currie and Almond (2011).

²For a general review of the policies, see Chzhen et al. (2019). For policies specific to young children, see UNICEF (2017).

³For a review focusing on policies, see Grimshaw and Rubery (2015).

⁴Notable examples of high-quality child care programs in the US are the Perry School Project or the Abecedarian Program. They are described and evaluated in, among others, Heckman et al. (2010b,a); Conti et al. (2016); Campbell et al. (2014). For a recent review on early childhood programs and their impacts, see Currie (2001) and Berlinski and Vera-Hernández (2019).

⁵Similar two-step games are implemented in Mazzeo (2002), Dranove et al. (2003), and Seim (2006). For an example of monopolistic competition among schools, see Urquiola and Verhoogen (2009).

instance, heterogenous managerial ability).

To see the relevance of our general equilibrium model, consider a policy that provides child care vouchers to all families. This would raise the demand for child care and, if the supply of child care centers did not change, it would raise the use and price of child care - as predicted by a partial equilibrium model. Yet our general equilibrium model allows for an endogenous adjustment of the supply, since the demand expansion would encourage the entry of new providers. As a result, we would predict a greater expansion of child care use, and a lower price increase, than in partial equilibrium. Particularly for the large-scale policies that have been often proposed to increase access to center-based care, the consideration of these general equilibrium effects is crucial.

We estimate the structural parameters of the model by Simulated Method of Moments using data from the Early Childhood Longitudinal Study (ECLS). We focus on the ECLS birth cohort data (ECLS-B), which is a nationally representative sample of children born in 2001, and followed from birth through kindergarten. We use the second wave, conducted between January 2003 and December 2003, since the children were two years old then. The data are well suited for this project because they contain not only measures of parental labor supply and child development, but also detailed information on child care arrangements and providers' quality. We focus on families with two-year old children since this is the age for which the private child-care market is prevalent and has a large potential for expansion.⁶

The estimated model fits the data well. In particular, it replicates the patterns of child care quality and price across families of different incomes, mothers' education, and mothers' labor force participation. Our parameter estimates show that the time spent both with parents and at child care centers has a positive impact on child development. Time spent with fathers has a weaker impact than time spent at child care centers; time spent with college-educated mothers has the strongest impact of all inputs.⁷ We also estimate complementarities in production such that, given the appropriate input mix, even college-educated mothers can work in the market without negatively impacting child development as long as they buy high-quality child care.

We use the estimated model to evaluate some policy interventions aimed at increasing access to center-based care. Some provide a demand subsidy, as is the case with vouchers for child care services (either universal, or conditional on family income or labor force participation) and universal cash transfers (which can be used to buy child care but also other goods and services). Others focus on the supply side, as is the case with regulation of child care quality, and public provision of child care. We compute the equilibrium of our model in the absence of policies ("baseline") and under each of these policies in order to study their impact on child development, child care use, child care markets, and mothers' labor force participation. We also provide a simple estimate of policies' fiscal cost.

We find that the greatest child development gains result from combining quality regulation and a voucher for families in which both parents work (i.e., a "work requirement voucher"),

⁶Enrollment rates in center-based care at ages 3 to 5 are relatively high in developed countries (Laughlin (2013) and Chzhen et al. (2019)), and provision is frequently public. The US has seen a stable increase in public pre-schooling programs targeting children ages three and four but not two.

⁷Throughout, "college" refers to four-year colleges. A "college-educated" parent has completed at least a four-year bachelor's program.

as this policy supports both the supply and demand of high-quality care. In this respect, a common concern is that quality regulation might raise not only quality but also prices, and thus render child care unaffordable to the low-income families whose children might benefit the most from it. In our simulations, prices do rise due to quality regulation yet not much because of the endogeneity of the supply side. In effect, by allowing only high-quality centers to operate, the policy encourages the entry of such firms, which must then compete on prices. Further, the limited price increase does not lower child care use because of the voucher, which actually expands use. Since the voucher is conditional on both parents working, the policy also increases mothers' labor force participation and slightly reduces the wage gender gap. As a result, this policy is the best option for a policymaker seeking to maximize mothers' labor force participation while ensuring high-quality child care at a low fiscal cost.

While this policy leads to a large expansion in child care use, the largest expansion comes from the public provision of child care. We model public provision as a "government" that chooses the combination of firms and qualities that operate in the market in order to maximize social surplus, and mandates that each firm's price be equal to its cost. By achieving the lowest possible average price, this policy yields the highest child care use of all policies. At the same time, the price becomes so low that some mothers no longer need to work in order to pay for child care, thus leading to a small decline in mothers' labor force participation. This policy does not maximize child development gains even though it raises the proportion of high-quality centers. The main reason is that, by providing high-quality care at a very low price, it crowds out parents' time with their children.

Although no other study models both the supply and demand of child care, numerous contributions have focused on only one side of the market. On the demand side, a large literature looks at parental labor supply, child care use, and child development. The first strand of this literature starts with the early contributions based on static labor supply models with unitary households that make employment and child care decisions.⁸ These papers use various US datasets, and estimate reduced form and structural models to identify the effect of child care costs on female labor supply and child care demand. They find that, consistent with the neoclassical model, maternal employment and the use of formal child care are sensitive to child care costs. Child care subsidies, particularly for the poor, can have important effects in increasing labor force participation and child care use. Our modeling of parents' labor supply is similar to that of these first-generation papers.

The second strand of the demand-side literature conducts impact evaluation of child care subsidy policies to estimate their effect on maternal employment and child development outcomes.⁹ For example, Baker et al. (2008) analyze the introduction of the \$5-dollar-a-day child care policy in Quebec, Canada. They find that this subsidy increases child care use and maternal employment, but has negative effects on child outcomes. Havnes and Mogstad (2011b) find that universal child care in Norway crowds out informal care arrangements without increases in maternal employment. The impact on long-run children's outcomes is heterogenous, with positive effects for those at the bottom of the income distribution

⁸See for example, Heckman (1974), Blau and Robins (1988), Ribar (1992), and Kimmel and Connelly (2007).

⁹There is also a growing literature on the child development impact of more generous parental leave policies (e.g., Dustmann and Schonberg (2012) and Carneiro et al. (2015)).

and negative at the top. Our model captures such heterogenous effects by introducing rich heterogeneity in both demand and supply, and by allowing some structural parameters to vary with mothers' education.

The third strand of the demand-side literature (e.g., Ruhm (2004), Bernal and Keane (2011), and Attanasio et al. (2015a)) studies the effect of maternal time as an input in the production of child development. The earlier papers estimate hybrid equations, where the parameters embody both technological properties of the production function and the effect of mediating variables. For example, Ruhm (2004) uses data on the children of individuals from the National Longitudinal Survey of Youth (NLSY), and finds that maternal employment during the first three years of life is negatively associated with children outcomes at ages three to five. Bernal and Keane (2011) evaluate the effect of child care versus maternal time using single mothers from the NLSY79. Exploiting exogenous variation in welfare policy rules, the estimates indicate a negative effect on child development. Finally, Attanasio et al. (2015a) estimate dynamic factor models of child cognitive and socio-emotional skills as a function of endogenous material and time investments, and find that time inputs are an important determinant of child outcomes.¹⁰ Our model includes both mother's and father's time as inputs in child development production, yet also includes time in child care centers, whose impact varies with center quality.

The fourth strand of the demand side literature employs a fully structural approach. Bernal (2008) estimates a dynamic model of employment, child care use, and production of child cognitive ability. This approach enables her to solve endogeneity issues related to maternal employment, and to perform counterfactuals. She estimates the model on a sample of single mothers from the NLSY79, and finds that a child care subsidy would increase overall utility for the parents but at the expense of child development. Also in the context of dynamic structural model of skill formation, Del Boca et al. (2014) use the Child Development Supplements of the Panel Study of Income Dynamics and exploit detailed data on how parents use time with their children, rather than the usual, coarse measure of maternal employment. Consistent with the rest of the literature, they find that parents' time inputs are relatively more important than material resources for children's cognitive development. Our paper also seeks to recover the structural parameters underpinning parents' decision. In addition, it models the supply side of the market. To handle the inherent complexity of endogenizing both demand and supply, it uses a static rather than a dynamic framework.

The supply-side literature is much more limited in number of papers and scope. Blau and Mocan (2002) estimate the relationship between cost and child care quality with data from a sample of US child care centers. Using a center quality measure similar to ours, they show that the supply of quality is moderately elastic with respect to prices and to the wages of child care workers. Based on US longitudinal data, Hotz and Xiao (2011) find that more stringent regulations on establishment inputs increase quality (particularly in high-income locations) but reduce child care provision (particularly in low-income locations). Bastos and Cristia (2012) develop a model of monopolistic competition to explain the behavior of child care providers in a large Brazilian city. Using Census data, they find a positive association between household income and child care quality, which is consistent with their model. Our

¹⁰Cunha et al. (2010) and Agostinelli and Wiswall (2016) also estimate dynamic models but do not separate returns to material and time investments.

paper also models the supply side as monopolistic competition. We build on the two-stage model of Mazzeo (2002), known in the Industrial Organization literature for being one of the first studies to estimate a model of entry with endogenous product characteristics (Reiss and Wolak (2007)). This framework has been extended and adapted to study competition among health maintenance organizations (Dranove et al. (2003)), in the video retail industry (Seim (2006)), and in the airline industry (Ciliberto and Tamer (2009)).

The remainder of the paper is organized as follows. Section 2 describes the model, and Section 3 presents the data. Section 4 discusses econometric issues, including identification and estimation procedure. Section 5 presents estimation results, and Section 6 reports policy simulations. Section 7 concludes.

2 Model

The economy is composed of two sets of agents: households – consisting of two parents and one child – and child care providers. In the household, parents are the decision makers and choose whether to work, and whether to send their child to an outside child care facility. Child care providers are private firms that supply outside child care services.¹¹ Potential providers decide whether to enter, and what quality and price to offer. In a unitary household framework, parents maximize household total utility, and the child care providers that operate in the market maximize profits. In equilibrium, no household wishes to alter its decisions on work or child care, and providers in the market make non-negative profits and do not wish to alter their price or quality.

2.1 Environment

2.1.1 Households

The economy is populated by a continuum of households of mass M . Each household includes two parents (mother, m , and father, f) and one child, and operates under a unitary model. There are J household types, each with measure M_j . Household types vary in labor market opportunities (wage offers), preferences, and an initial level of child development. In what follows we define the environment conditioning on a given type j and mostly suppress the dependence on j to facilitate exposition.

There are $(T + 1)$ possible child care arrangements, as the household can either choose one of the T child care providers available in the economy, or can care for the child at home. A household’s utility function under a given child care arrangement, t , is given by:

$$U_t = c^{\alpha_c} q^{\alpha_q} l_f^{\alpha_f} l_m^{\alpha_m} \exp(\varepsilon_t), \quad (1)$$

where c is total household consumption, q is child development, l_f and l_m denote father’s and mother’s leisure, respectively, and ε_t represents the household’s idiosyncratic preferences for child care arrangement t . Note that each household has a $(T + 1)$ vector of idiosyncratic

¹¹We use the terms “child care provider” or “outside care” for non-parental, center-based child care.

preferences for child care arrangements, ε . These vectors vary among households of a given type, and across types as well.

We follow Del Boca et al. (2014) by imposing homogeneity of degree one in our Cobb-Douglas formulation, or $\alpha_c + \alpha_q + \alpha_f + \alpha_m = 1$. The household faces a budget constraint:

$$c = h_f w_f + h_m w_m - t_d p + I, \quad (2)$$

and a set of time constraints:

$$\text{Father : } h_f + t_f + l_f = 16 \quad (3)$$

$$\text{Mother : } h_m + t_m + l_m = 16 \quad (4)$$

$$\text{Child : } t_f + t_m + t_d = 16 \quad (5)$$

where $\{w_f, w_m\}$ are parents' hourly wage offers (which, again, vary across household types) and $\{h_f, h_m\}$ their labor supply measured in hours per day; t_d denotes the daily hours spent by the child in outside care at an hourly price of p ; I is non-labor income, and $\{t_f, t_m\}$ are the daily hours of care provided by the parents.

The interpretation of the above is as follows. We assume that eight hours a day are devoted to sleeping, and therefore parents and children are endowed with 16 hours a day each. Parent x can devote her time to working h_x hours, caring for her child for t_x hours, or enjoying l_x hours of leisure. The child spends t_f and t_m hours a day with father and mother, respectively, and t_d hours a day in outside child care. Since the model is static, all income is spent either in consumption or in child care services.

Child development is produced according to a Constant Elasticity of Substitution (CES) production function, in which the current inputs are the time of parents, $\{t_f, t_m\}$, and the time spent with outside care providers, t_d . In a value-added sense, these are added to the *ex ante* child development, q_0 , and produce (current) child development as follows:¹²

$$q = (\gamma_0 q_0^r + \gamma_f t_f^r + \gamma_m t_m^r + \gamma_k t_d^r)^{\frac{1}{r}} \quad (6)$$

where r is the parameter governing the substitution elasticity, equal to $1/(1 - r)$. This functional form allows both for substitutability and complementarity among inputs; it nests other common functional forms such as linear ($r = 1$) and Cobb-Douglas ($r = 0$). The specification also allows mother's productivity to differ from father's as $\{\gamma_f, \gamma_m\}$ are the returns to the time spent with the father and mother, respectively. The return to time spent at the child care center is γ_k , which can be viewed as the center's quality. The *ex ante* child development, q_0 , captures the contribution of all past inputs to child development, in the spirit of the child development process modeled by Cunha et al. (2010)

Household types reflect the *ex ante* heterogeneity of households, prior to their decisions on labor supply or child care. Each type j is characterized by a vector of preference parameters over consumption, child development, and leisure, $\{\alpha_c(j), \alpha_q(j), \alpha_f(j), \alpha_m(j)\}$; a vector of wage offers for the two parents, $\{w_f(j), w_m(j)\}$; and an initial child development, $\{q_0(j)\}$. These types reflect non-idiosyncratic heterogeneity. At the same time, households *within a*

¹²This functional form is common in the literature even when the inputs considered are different from ours and do not include child care providers. The seminal example is Cunha et al. (2010).

type may differ in idiosyncratic preferences for a given child care arrangement t . Idiosyncratic preferences are crucial to generate horizontal differentiation in the child care market, because they create a taste for variety on the part of households. For example, two households with the same wage offers, preferences, and initial child development may have different preferences over unobserved providers' attributes such as geographic location. As a result, these households may choose centers of similar price and quality yet different locations. In addition to horizontal differentiation, the model allows for vertical (quality) differentiation, as parents value child care quality but differ in their ability to pay for it.

2.1.2 Child care providers

The economy includes K potential child care providers that may choose to enter the market after paying a fixed cost F . We assume that providers are profit-maximizers and can only open one center each. Centers can, in principle, differ in quality, γ , and price, p . For computational reasons we assume only two qualities, high and low (γ^H and γ^L respectively). Potential entrant k has a type given by its costs of providing low and high-quality care, $\{c_k^L, c_k^H\}$, which are drawn from a joint distribution with $E[c_k^L] \leq E[c_k^H]$. These draws are common knowledge.

Given its costs, center k chooses whether to provide γ^H or γ^L . Since providers have different costs, two providers may offer the same quality yet incur different costs - due, perhaps, to different managerial abilities. Thus, center k 's variable cost for providing quality γ_k is given by:

$$c_k(\gamma_k) = \begin{cases} c_k^L, & \text{if } \gamma_k = \gamma^L \\ c_k^H, & \text{if } \gamma_k = \gamma^H \end{cases} \quad (7)$$

Following Mazzeo (2002), potential entrants engage in a two-stage game. The first stage is a Stackelberg game in which firms choose among the following options: enter the market offering low quality, enter offering high quality, and not enter at all. As a result, the first stage determines the market configuration \mathfrak{S} , which is the set of T providers active in the market as well as their qualities. In the second stage, the active providers engage in a Bertrand game of simultaneous competition by choosing prices, p , conditional on \mathfrak{S} and on households' demand for child care services.

2.2 Equilibrium

2.2.1 Household decision problem

In a household, parents make labor supply and child care choices. They decide how to allocate their time between work and child care; how many hours (if any) to purchase from a child care center, and which center (if any) to purchase from. For tractability and consistency with our empirical application, we apply three restrictions to the set of possible choices. First, labor supply can only be full-time, equal to eight hours a day. Second, a child attending a child care center must spend eight hours a day there. Third, the father always supplies

labor.¹³ As a result of these restrictions, households in which both parents work must always hire child care because we assume that the child can never be left without adult supervision (equation (5)). In other words, parents can choose whether both work or not, and whether to use child care or not, but cannot choose how many hours to work or how many hours to send the child to an outside provider.

In contrast, parents can allocate their non-work time freely between leisure and at-home child care. While child care *must* be purchased by households with two working parents, it *may* also be purchased by households in which only father works. Importantly, parents must still choose hours of leisure even if they purchase child care. For instance, if mother stays home with the child and hence provides eight hours of child care while father works, she still has eight hours of non-sleep time to allocate between additional child care, t_m , and leisure, l_m . Note that while h_m and h_f are discrete, t_m and t_f are continuous. Further, note that a parent who works and sends the child to child care can still choose to spend some of his/her non-work time in child care, during the hours that he/she overlaps with the child at home.

The overall optimization problem consists of maximizing the household utility in (1) with respect to $\{c, q, l_f, l_m\}$ under the constraints (2)–(6) and our choice set restrictions, $h_f = 8$, $h_m \in \{0, 8\}$, $t_d \in \{0, 8\}$, and $t_d = 8$ if $h_m = 8$. Recall there are T child care providers, each one offering a potentially different combination of price and quality, and let k denote a specific child care center. Let $k = 0$ represent the option of home-based care, provided by mother (with $h_m = 0$ and $t_d = 0$). Overall, the household faces many possible $\{h_m, t_d, k\}$ options.

We can view the household utility maximization as a two-step decision. In the first, households choose mother’s optimal child care time, t_m , from a continuous choice set for each possible $\{h_m, t_d, k\}$. Given (5), maximizing over t_m is equivalent to maximizing over t_f . For a given $\{h_m, t_d, k\}$, the optimal t_f is

$$t_f^* = t_f^*(h_m, t_d, k) \quad (8)$$

Let ω denote a generic (h_m, t_d, k) combination, and let Ω be the set of all such combinations. Then the solution to the second step (and the household’s optimal decision) consists of the vector $\omega^* \in \Omega$ such that:

$$u(t_f^*(\omega^*)) \geq u(t_f'(\omega')), \quad \forall t_f' \in [0, 16] \text{ and } \forall \omega' \in \Omega \quad (9)$$

We assume that ε_t follows a type I extreme value distribution with scale $\mu_\varepsilon > 0$. We also assume that it is independently and identically distributed across the $(T + 1)$ child care arrangements available to a given household, and across households. For the option of home-based care ($k = 0$), $p_0 = 0$ and γ_0 depends on mother’s education, as explained below. Hence, the probability that a household of type j choose work hours h_m and choose child care option k , with quality γ_k and price p_k , is equal to

$$P_{jk}(\mathbf{p}, \boldsymbol{\gamma} \mid \omega_k) = \frac{\exp(\tilde{U}_{jk}(p_k, \gamma_k; \omega_k) / \mu_\varepsilon)}{\sum_{r=0}^K \exp(\tilde{U}_{jr}(p_r, \gamma_r; \omega_r) / \mu_\varepsilon)}, \quad (10)$$

¹³We have restricted our sample only to households in which the father supplies labor in the market.

where

$$\tilde{U}_{jk}(p_k, \gamma_k; \cdot) \equiv \alpha_c \ln(c(p_k)) + \alpha_q \ln(q(\gamma_k)) + \alpha_f \ln(l_f) + \alpha_m \ln(l_m), \quad (11)$$

$c(p_k) = 8w_f + h_m w_m - p_k + I$, and $q(\gamma_k) = (\gamma_0 q_0^r + \gamma_f t_f^r + \gamma_m t_m^r + \gamma_k \mathcal{S}^r)^{\frac{1}{r}}$. This probability can be seen as the share of type j households that make this particular choice. The resulting market demand for child care center k , with quality γ_k and tuition p_k , is equal to

$$N_k(\mathbf{p}, \gamma) = \sum_{j=1}^J \left[\int_{\omega \in \Omega_j} P_{jk}(\mathbf{p}, \gamma | \omega) \right] M_j. \quad (12)$$

2.2.2 Providers' game

Following Mazzeo (2002) we assume that firms' optimal decisions are the solution of a two-stage game. The first stage is the *entry stage*, in which potential entrants decide whether to enter the market as a low-quality provider (offering γ^L), as a high-quality provider (offering γ^H), or not to enter at all. The second is the *competition stage*, in which actual entrants compete by setting prices given the qualities chosen by all firms at the entry stage.

The entry stage consists of a Stackelberg, sequential sub-game. Potential entrants make their decisions sequentially, in an order given by their expected profitability or efficiency. Appendix C provides details on the order of entry. Intuitively, the most efficient firm decides on entry and quality knowing that subsequent entrants will take its decision into account. The second most efficient firm makes its decisions taking into account the most efficient firm's choices, and knowing that subsequent entrants will take into account the decisions made by the first and second entrant. In general, entrant k makes its entry and quality decision observing the choices made by the previous entrants, $k' < k$, and knowing that the upcoming entrants, $k'' > k$, will take its decision into account. If potential entrant k decides to enter, it chooses the quality that maximizes profits at the second (competition) stage given the decisions made by the previous entrants. By establishing which firms enter the market and their qualities, the first stage determines the equilibrium market configuration, \mathfrak{S}^* .

Since players have complete information about their rivals' types and chosen qualities, provider k knows the realized cost vector (c_v^L, c_v^H) for every $v \neq k$. However, providers are uncertain about each consumer's ε and only have information about its distribution. Therefore, the competition stage is a Bertrand game in which firms simultaneously choose prices in order to maximize expected profits given \mathfrak{S}^* . The problem of firm $k \in \mathfrak{S}^*$ is

$$\max_{p_k} E(\pi_k | \mathfrak{S}^*) = [p_k - c_k(\gamma_k)] N((p_k, \mathbf{p}_{-k}), \mathfrak{S}^*) - F, \quad (13)$$

where p_k is the price charged by provider k and \mathbf{p}_{-k} represents the prices charged by k 's competitors.

2.2.3 Equilibrium Definition

Bringing together the demand and supply for child care defines the market equilibrium. Households choose labor supply, leisure, child care time, child development, and consumption

by solving the utility maximization problem defined in Section 2.2.1. This leads to the demands defined by equation (12).

On the supply side, potential entrants choose whether to enter the market and with what quality; conditional on entering, they choose their price. In the entry stage, the sub-game perfect Nash equilibrium is obtained by backward induction as in a standard Stackelberg game, such that only centers with expected positive profits operate in the market. In the competition stage, the Nash equilibrium consists of a vector of prices \mathbf{p}^* that satisfies the first-order condition of profit maximization for every k :

$$p_k^* = c_k(\gamma_k) - N((p_k^*, \mathbf{p}_{-k}), \mathfrak{S}^*) \left[\frac{\partial}{\partial p_k} N((p_k^*, \mathbf{p}_{-k}), \mathfrak{S}^*) \right]^{-1} \quad (14)$$

Mazzeo (2002) demonstrates that this two-stage game has a unique equilibrium for a given number of firms.

Definition 1 *Equilibrium Definition.*

An equilibrium in this model specifies the set of child care providers available to households as well as their prices and qualities, and the decisions made by households with respect to work and child care, such that:

1. *the child care providers in parents' choice sets are indeed available in the market;*
2. *every child care provider in the market makes non-negative profits;*
3. *neither parents nor child care providers wish to alter their decisions.*

Due to the discreteness of household and firm choices, as well as the finite number of household types and potential firms, the equilibrium does not have a closed-form solution and must be solved numerically for a given set of parameter values. Appendix A describes the quantitative version of our model, and Appendices B and C describe the equilibrium computation.

3 Data

3.1 Sources

To estimate the model, we use data from the US Early Childhood Longitudinal Study birth cohort (ECLS-B). This study follows individuals from birth through kindergarten for a nationally representative sample of children born in 2001, and includes several waves in order to capture children's development as they grow. We use the second wave, collected between January and December 2003, when the children are two years old. In this wave the ECLS-B uses instruments to assess child development in the physical, cognitive, and socio-emotional domains. We focus on the cognitive domain, measured with the Bayley-Short Form instrument (Research Edition-Mental), which contains measures of general cognitive ability such as problem solving and language acquisition.

The ECLS-B also collects household demographic data (e.g., household members, parental education, parental age, marital status, and living arrangements) and parental labor market information (e.g., labor force participation, hours of work, and hourly wages). Moreover, the ECLS-B contains detailed information on child care arrangements (e.g., parental care, relative care, non-relative care, center-based care), hours of use for each one, and hourly child care expenses. Further, it provides an assessment of the quality of center-based care using an instrument designed by developmental psychologists, the ITERS (Infant/Toddler Environment Rating Scale).¹⁴

Starting from the full dataset of the second wave, we arrive at our estimation sample by imposing three main restrictions: we focus on two-parent families with only one child, who is two years old; we focus on families in which the father works full time, and in which the mother either works full time or does not work; and we consider only two child care arrangements, namely parental care and child care center. The first restriction greatly reduces the dimensionality of the problem by making parents choose child care for only one child. It also avoids the modeling of marriage, fertility, and economies of scale in child care arrangements. The first and second restrictions together lead to less variation among the choices made by families, and reduce the need to explicitly model intensive margin decisions (related, for instance, to the number of hours worked). The third restriction is the most costly when studying child care markets because it rules out some popular child-care arrangements such as nannies or relatives. However, since we want to model the endogenous supply of child care centers, introducing nannies or the use of relatives would require solving a “boundary of the firm” problem (Hart, 1995) which is beyond our scope.¹⁵ As a result of these restrictions, our estimation sample is composed of families in which, on average, parents are more educated than in the full dataset and earn higher wages, and in which mothers have a lower labor force participation rate. Appendix D provides further comparison between our estimation sample and the full dataset.

3.2 Descriptive Statistics

Our estimation sample consists of 450 households that meet the criteria described above.¹⁶ In Table 1 we define the main variables used in estimation and report descriptive statistics. About 30 percent of families use a child care center as the main form of care and pay on average 4.6 dollars per hour of full-time care. The average quality of child care centers, as measured by the ITERS score, is equal to 4.167 and hence above satisfactory, as values of

¹⁴ The ITERS scale contains 39 items organized in seven sub-scales: space and furnishings, personal care routines, listening and talking, activities, interaction, program structure, and parents and staff. We focus on the overall score, which is the mean over the seven sub-scales. This score is not available for all daycare centers surveyed but only for those that accepted to be observed (approximately 45 percent of all centers in the sample).

¹⁵ Modeling the use of nannies or relatives is akin to a boundary-of-the-firm problem because parents can be viewed as firms deciding whether to produce the good at home by themselves, whether to hire somebody (e.g., a nanny) who will do it for them with inputs they will provide, or whether to outsource the production altogether to a childcare center.

¹⁶ The sample size is rounded to the nearest 50 in order to comply with ECLS-B confidentiality rules.

4 and 5 denote satisfactory and good quality, respectively.¹⁷ This average is in line with previous work for the US population (e.g., Blau and Mocan (2002)), and is the same as the mean value of ITERS for children aged 2 in the full ECLS sample. The average child cognitive score is approximately equal to 128, which we divide it by 10 to place all variables in a similar scale for estimation.

In terms of parental labor market outcomes, we observe the usual gender differentials, as women participate significantly less than men and earn, on average, a lower wage. The gender differential in participation rates is much more pronounced than in a representative sample of the population, as only 30 percent of women work full-time.¹⁸ For similar reasons, the average gender wage gap (of about 10 percent) is smaller than in the overall population. Finally, about 50 percent of the parents in our sample have completed at least college, which yields a sample of households with relatively high educational attainment.

4 Econometric Issues

4.1 Empirical Specification

Before discussing identification and estimation, we describe issues related to household heterogeneity and functional form assumptions from our empirical specification.

4.1.1 Household heterogeneity

From Section 2, recall that household types vary in wage offers, preferences, and initial child development. When computing the equilibrium for our estimation, we treat each of the 450 households in our sample as a different type. To provide a better fit of the data and enrich our policy analysis, we add an additional source of *ex ante* household heterogeneity, namely parents’ educational attainment (“education,” for short). Hence, for each observed household we take from the data its initial child development, q_0 , and parents’ education. We classify the latter into high school and college, respectively. We use the term “college mother” for mothers who have completed a four-year bachelor’s or more, and “high school mother” for all others. Thus, in our empirical specification household types differ in four dimensions: parents’ education, wage offers, preferences, and initial child development. Parental education affects some model parameters referring to wage offers and child development production, as described below.

As for q_0 , it serves as the initial condition in the value-added specification of child development production in equation (6). We measure it as the child’s cognitive score at age one. See Appendix A for further details on the simulation setup.

¹⁷ITERS scores range from 1 to 7, with 1= inadequate quality and 7=excellent quality. We report the average over all 7 sub-scales, as explained in footnote 14.

¹⁸This is mainly due to our focus on households where the father works full-time. See Table D.1.

4.1.2 Functional forms

We make five sets of assumptions. First, we address the heterogeneity in household preferences. Recall that household type j has a vector of preference parameters over consumption, child development, and leisure: $\boldsymbol{\alpha}(j) \equiv \{\alpha_c(j), \alpha_q(j), \alpha_m(j), \alpha_f(j)\}$. As in Del Boca et al. (2014), we assume that these utility parameters are i.i.d. draws from a distribution constructed as follows. We define the 3×1 normally distributed random variable $\mathbf{v} \sim N(\boldsymbol{\mu}_v, \boldsymbol{\Sigma}_v)$, where:

$$\boldsymbol{\mu}_v = \begin{bmatrix} \mu_{v1} \\ \mu_{v2} \\ \mu_{v3} \end{bmatrix} \quad \text{and} \quad \boldsymbol{\Sigma}_v = \begin{bmatrix} \sigma_{v1}^2 & 0 & 0 \\ 0 & \sigma_{v2}^2 & 0 \\ 0 & 0 & \sigma_{v3}^2 \end{bmatrix} \quad (15)$$

For household type j , we draw a $\mathbf{v} = [v_1, v_2, v_3]$ vector and generate j 's utility parameters:

$$\alpha_c(j) = \frac{\exp(v_1)}{1 + \sum_{r=1}^3 \exp(v_r)} \quad (16)$$

$$\alpha_q(j) = \frac{\exp(v_2)}{1 + \sum_{r=1}^3 \exp(v_r)} \quad (17)$$

$$\alpha_f(j) = \frac{\exp(v_3)}{1 + \sum_{r=1}^3 \exp(v_r)} \quad (18)$$

$$\alpha_m(j) = \frac{1}{1 + \sum_{r=1}^3 \exp(v_r)} \quad (19)$$

This procedure is flexible and uses a convenient distributional assumption. It also satisfies the theoretical constraints over the four-dimensional vector $\boldsymbol{\alpha}(j)$ because the sum of the four scalars is equal to one.

Second, we address heterogeneity in household labor market opportunities. Recall that household type j faces wage offers $\{w_f(j), w_m(j)\}$. We assume wage offers are drawn from a joint lognormal distribution, frequently used in the literature for its good fit. We allow the distribution to depend on parents' education, and also to capture possible assortative (on education) mating between parents. Specifically, the parental log-wage offers follow a normal distribution with mean

$$\boldsymbol{\mu}_w[d_f^c(j), d_m^c(j)] = \begin{bmatrix} \mu_{w_f} + d_f^c(j)\mu'_{w_f} \\ \mu_{w_m} + d_m^c(j)\mu'_{w_m} \end{bmatrix} \quad (20)$$

and variance-covariance matrix:

$$\boldsymbol{\Sigma}_w[d_f^c(j), d_m^c(j), d^s(j)] = \begin{bmatrix} \sigma_{w_f}^2 + d_f^c(j)(\sigma'_{w_f})^2 & (\rho + d^s(j)\rho')\sigma_{w_f}\sigma_{w_m} \\ (\rho + d^s(j)\rho')\sigma_{w_f}\sigma_{w_m} & \sigma_{w_m}^2 + d_m^c(j)(\sigma'_{w_m})^2 \end{bmatrix} \quad (21)$$

where $d_f^c(j)$ and $d_m^c(j)$ are equal to 1 if father and mother have completed college, respectively, and $d^s(j) = 1$ if both parents have the same education (e.g., if both completed either high school or college). Parameters μ'_w and σ'_w capture possible shifts in the mean and variance of wage offers for college relative to high school graduates. Parameter ρ reflects the intensity of assortative mating in wage offers among parents of different educations; the intensity changes to $(\rho + \rho')$ when both have the same education.

Third, we set the value of the non-labor income, I . We find the minimum level of daily household expenditures from our sample, and set I equal to the negative of this level. This is needed to ensure that household income, net of child care expenses, covers at least a minimum consumption level (for instance, in housing and food).¹⁹

Fourth, we allow mother’s education to play a role in the production of child development. Recall from (6) that her productivity in child development production is γ_m . We set:

$$\gamma_m[d_m^c(j)] = \gamma_m + d_m^c(j)\gamma'_m \quad (22)$$

so that college completion can potentially shift the returns to mother’s time. We limit the effect of education to mothers due to data limitations and given previous findings on the asymmetric roles of mothers and fathers in child-rearing.²⁰

Fifth, we assume a parametric distribution for the cost of providing child care services. Recall from (7) that each potential child care provider k has its own variable cost of providing low-quality and high-quality care, (c_k^L and c_k^H , respectively). We assume that c_k^L and c_k^H belong to a lognormal distribution with parameters $\{\mu_c^L, \sigma_c^L\}$ and $\{\mu_c^H, \sigma_c^H\}$, respectively. In addition, all firms face a fixed cost, F , which is independent of their quality.

4.2 Identification

The model is characterized by five sets of parameters, related to the following:

1. child development production function;
2. child care center costs;
3. utility function;
4. wage offer distributions;
5. idiosyncratic preferences with respect to child care arrangements.

We start by discussing the identification of the child development production function defined in equation (6). The function includes five parameters: parents’ productivity γ_f and γ_m , child care center’s productivity γ_k , the impact of the ex-ante child development γ_0 , and

¹⁹In our empirical specification, I is equal to -\$40 per day. We have found that $I \geq 0$ leads to unrealistically high levels of child care spending and hence too little consumption. By limiting the disposable income available for child care, a negative I predicts realistic values of child care spending.

²⁰Correlation studies indicate that “the productivity of women in home production appears to increase as their education increases” (Schultz (1993), p 53) while fathers’ education seems to be significantly less important. Admittedly, the conclusion is more nuanced among the several studies that use arguably exogenous variation on maternal education to look at its effect on children in the short and long run (for a survey, see Black and Devereux (2010)). An example is Carneiro et al. (2013) finding that maternal education has a positive effect on children when they are 7 to 8 years old but a negative one when they are younger than 24 months old. In addition, the reduced form relationship between mother’s education and children outcomes estimated by these studies involves a variety of mechanisms such as labor market participation and assortative mating in the marriage market. Clearly, at least the second one is not independent of father’s education.

the substitution parameter r . Given the model, the observed data allows us to identify these parameters as follows. We observe the time input from the child care center (t_d) because the data reports child care use and we assume that parents can only use centers for full-time care. Although we do not directly observe the parents’ optimal time inputs (defined in equation (8)), these are closely related to their observed labor market participation and child care use through their time constraints (equations (3)–(5)). We observe child development since we assume that the Bayley child development outcome indicator measures it without error. Thus, the observed Bayley indicators at ages one and two deliver the ex-ante and ex-post level of development, q_0 and q respectively. We have less information on child care center quality, but we observe centers’ ITERS measure for about half of the centers in the sample (which enroll about half of the children).²¹ Since we have assumed that γ_k can assume only two values, we classify centers with ITERS > 4 as high-quality, and all others as low-quality.²² Given q , q_0 and γ_k , the variation in hours worked by parents, within and across educational levels, identifies γ_f , γ_m , γ'_m and r .

As described in Section 4.1, child care center costs are characterized by five parameters: those from the lognormal distributions for the variable costs, $\{\mu_c^L, \sigma_c^L\}$ and $\{\mu_c^H, \sigma_c^H\}$, and the fixed cost F . We do not estimate F due to data limitations. Instead, we set it equal to \$204.56 per day based on industry cost reports (Mitchell and Stoney, 2012).²³ Regarding variable costs, we face the usual selection problem in that we do not observe all potential entrants but only those that do enter - and not even for these we do observe variable costs. Identification is therefore based on two pieces of information: price charged and quality provided. We observe prices under the assumption that centers do not discriminate prices among their users. We observe quality in the aggregate fashion described above. Given F , and the observed market share for high- and low-quality centers together with the price distribution for each quality level, we identify the variable cost parameters given that active firms make a non-negative profit.

Given our assumed heterogeneity in utility function parameters, we must identify six parameters: $\{\mu_{v1}, \mu_{v2}, \mu_{v3}\}$ and $\{\sigma_{v1}, \sigma_{v2}, \sigma_{v3}\}$. Household preferences affect trade-offs among leisure, consumption, and child development. As described in Section 3, we observe wages,

²¹Looking at the variables of interest, we did not detect any particular selection on families with or without ITERS reported. Therefore, we assume that the observations with a reported ITERS measure constitute a representative sample in the given cell for which we compute the moments of the SMM estimation procedure (explained in Section 4.3).

²²ITERS scores range from 1 to 7. ITERS = 4 means ‘satisfactory quality’ and = 5 means ‘good quality’ (see Section 3.1 for more details). Therefore, a threshold of 4 seems a natural choice. Still, it remains an arbitrary choice over which we have performed some robustness checks. The difference in estimation results is negligible when we use a threshold of 5, but it grows with thresholds of 6 or higher since the proportion of high quality center decreases substantially.

²³ We use data from the Generic Cost Model for child care centers provided in Mitchell and Stoney (2012). The document is compiled by the *Alliance for Early Childhood Finance* with the objective of modeling “budgets for a proposed center and to understand the costs of operating a better quality center”. In the fixed cost calculation we include the non-personnel costs for the building (rent/lease, insurance, utilities, maintenance/repair/cleaning) and the setup costs (audits, fees and permits, telephone and internet). We compute all costs at 2003 prices and assume that the typical center size is two classrooms. A value of about \$200 a day is comparable with the fixed costs resulting from occupancy and overhead reported by Helburn and Howes (1996). We renormalize the per-child measures of the Generic Cost Model by assuming an average size of two classrooms with 14 children.

labor market participation, and child care use for each household. The choices made by the “average” household identify “average” preferences and hence $\{\mu_{v1}, \mu_{v2}, \mu_{v3}\}$, whereas the variation in choices across household types identifies the variance in preferences and hence $\{\sigma_{v1}, \sigma_{v2}, \sigma_{v3}\}$.

Under the functional form assumptions proposed in Section 4.1, the joint wage offer distribution is characterized by ten parameters. Each gender and education group is allowed to have different location and scale parameters, leading to eight parameters:

$$\{\mu_{w_f}, \sigma_{w_f}, \mu'_{w_f}, \sigma'_{w_f}, \mu_{w_m}, \sigma_{w_m}, \mu'_{w_m}, \sigma'_{w_m}\}. \quad (23)$$

In addition, we allow for within-family correlation between mother’s and father’s offer, and allow the correlation to vary depending on whether mother and father have the same or different educational levels. This adds two parameters, ρ and ρ' . As is common, we do not observe the distribution of the wage offers but only the distribution of accepted wages. However, the latter is a truncation of the former given the model and our distributional assumptions on wage offers. Since we assume a distribution that can be recovered from its truncation and can characterize the parent’s decision rule leading to the truncation, we identify wage offer parameters based on observed hourly wages.²⁴ We identify parameters that are gender- and education-specific based on observed mothers’ and fathers’ education.

Finally, recall that the idiosyncratic preference with respect to child care arrangements is distributed type I extreme value with scale parameter μ_ε . If all households of a given type had the same idiosyncratic preferences, they would all choose the same child care arrangement. Thus, the greater the heterogeneity in idiosyncratic preferences, the greater the horizontal differentiation in the child care market. The dispersion of households’ observed child care choices regarding use, price, and quality conditional on household education, income, and child development is therefore the main source of identification for μ_ε .

4.3 Estimation Method

We estimate the model by Simulated Method of Moments (SMM). This method is particularly convenient in our context as we lack a closed-form solution for the equilibrium. Let $\boldsymbol{\theta}$ be the column vector with the $P = 28$ parameters to estimate. Our estimator, $\hat{\boldsymbol{\theta}}_{\text{SMM}}$, is such that:

$$\hat{\boldsymbol{\theta}}_{\text{SMM}} = \arg_{\boldsymbol{\theta}} \min [\mathbf{s}(\boldsymbol{\theta}) - \mathbf{z}]' \mathbf{W}^{-1} [\mathbf{s}(\boldsymbol{\theta}) - \mathbf{z}] \quad (24)$$

where \mathbf{z} is the column vector of B sample moments that we intend to match, and $\mathbf{s}(\boldsymbol{\theta})$ is the column vector of the corresponding moments simulated at parameters $\boldsymbol{\theta}$. The weighting matrix, \mathbf{W} , is diagonal; the diagonal consists of the bootstrapped standard deviation of the B sample moments.²⁵ We perform full-solution estimation, which requires the equilibrium computation for every value of $\boldsymbol{\theta}$.

We choose the $B = 36$ moments to match in the estimation based on the identification strategy described in Section 4.2. We report their sample value in the columns labeled

²⁴The lognormal is a recoverable distribution, i.e. it can be identified by observing its truncation at a known truncation point (Flinn and Heckman, 1982).

²⁵The weighting matrix is useful to give all moments a similar scale, and to reduce the estimation’s sensitivity to imprecise sample moments.

“Sample” in Table 4. We match moments related to child care markets, child development, and parents’ labor market. We compute these moments conditional on mother’s education and labor force participation status. For child care markets, we match average child care use, price average and standard deviation, and proportion of children in high-quality centers. For child development, we match average and standard deviation of the Bailey child development indicator. For parents’ labor market, for each parent we match the average and standard deviation of wages. In addition, we match mother’s labor force participation status by mother’s education. While we would prefer to condition our moments both on mother’s and father’s education given that the labor market parameters are gender- and education-specific, a substantial number of (father, mother) education combinations have very few observations.²⁶ Thus, we condition only on mother’s education, and rely on married couples’ assortative mating in education in order to estimate education-specific parameters for fathers.

5 Estimation Results

Table 2 reports the point estimates for the model’s structural parameters. Panel a corresponds to the household utility function (equation (1)). The first six parameters characterize the distribution of preferences over consumption, leisure, and child development; the seventh parameter characterizes idiosyncratic preferences for child care arrangements. Given the parameterization described in Section 4.1, it is more convenient to discuss preference heterogeneity by looking at Table 3’s panel a. This reports the mean and standard deviations of the utility weights $\alpha(j)$ and shows that households put, on average, a similar weight on consumption and child development, yet a much lower weight on parents’ leisure. Further, the average value of mother’s leisure is about twice as high as that of father’s.

Our estimates are in broad agreement with others in the literature. Using the same preference specification as ours, Del Boca et al. (2014) obtain estimates for the average utility weights on consumption and child development that are similar to ours.²⁷ They estimate that father’s and mother’s time have similar utility weights, while we estimate a higher one for mother’s. Our finding reflects, at least partly, the greater (positive) gap between father’s and mother’s labor supply in our sample. Our estimated dispersion in utility weights is lower than that in Del Boca et al. (2014) yet still comparable. Our estimate of μ_ε suggests that the role of idiosyncratic preferences is substantive yet smaller than that of non-idiosyncratic preferences.²⁸ This gives rise to a substantive amount of horizontal differentiation in the child care market, although vertical differentiation prevails and households sort among high- and low-quality providers based on their willingness to pay.

²⁶For example, we only observe eight households where the mother has completed high school, the father has completed college, and the mother works. Similarly, we only observe 27 households where the father has completed high school, the mother has completed college, and she does not work.

²⁷We compare our estimates with their estimates for the one-child case, reported in Column 1 of Table 5. They estimate average weights on consumption and child development equal to 0.257 and 0.353 respectively, whereas ours are 0.360 and 0.374.

²⁸As equation (10) shows, the size of μ_ε relative to that of the utility function parameters determines the relative role of the idiosyncratic preferences. By construction, utility function coefficients add up to 1. Since the estimate for μ_ε is equal to $0.816 < 1$, idiosyncratic preferences are important yet not as much as non-idiosyncratic preferences.

Table 2’s panel b shows parameter estimates for the child development production function (6). The production function parameters include the technological parameters of each input, γ_0 through γ_H , and the parameter governing the elasticity of substitution, r . The technological parameter estimates show a strong positive impact for past child development, mother’s time, and time spent at high-quality centers, yet a substantially weaker impact for father’s time. Moreover, high school mothers have a substantially weaker impact than college mothers, and a slightly weaker impact than high-quality centers. College mothers have the strongest impact of all inputs.

Our ranking of the productivity of mother’s and father’s time is consistent with previous literature.²⁹ However, the comparison between parents’ time and child care center time is less common in the literature and has interesting implications. Since the point estimate of r is smaller than one, the elasticity of substitution among inputs is larger than one (equal to 3.61). This implies a degree of complementarity among inputs, and non-linearity in their marginal returns. For example, even though college mothers’ time is the most productive input, under certain conditions it is possible to substitute mother’s time with another input and still raise child development. Consider the average household with a college mother. If mother does not work and the child does not attend child care, average child development is approximately equal to 11.9. But if mother works and the child attends a high-quality child care center, then child development is about 4 percent higher. Various degrees of complementarity, particularly between the previous period’s child development and other inputs, have also been found in the literature.³⁰

Table 2’s panel c shows wage offer parameters, which are easier to interpret by looking at their implied moments. Table 3’s panel b reports the mean and standard deviation of wage offers, overall and by parents’ education. Results show the expected ranking: higher wage offers for college graduates than high school graduates; higher wage offers to men than women. The wage offers are particularly low, on average, for high school mothers. The correlation coefficients ρ and ρ' indicate significant assortative mating in wage offers, leading to a relatively low wage offer gender gap (of about 20 percent) for couples where both parents are college graduates.

Table 2’s panel d shows the supply-side cost parameters. As with wage offers, these are easier to interpret by looking at the implied mean and standard deviation of variable costs reported in Table 3’s panel c. As expected, variable costs for high-quality centers have a higher mean than those for low-quality centers (about 50 percent higher). They also have a higher variance to accommodate some relatively high prices observed in the data.

In Table 4 we report the value of the observed (“sample”) and predicted (“simulated”) moments; the latter are evaluated at our parameter estimates. Our model fits the data reasonably well, as predicted moments rarely differ from observed moments by more than 5 or 10 percent. Importantly, moments of child development are well matched, and we reproduce

²⁹Del Boca et al. (2014) distinguish between active and passive time with the child and find a large difference in the impact of mothers’ and fathers’ active time. We model *total* time, and estimate a larger difference in the impact of mothers’ and father’s (total) time.

³⁰Cunha et al. (2010), Agostinelli and Wiswall (2016), Attanasio et al. (2015a), and Attanasio et al. (2015b) find complementarities in production. Their models are dynamic and include different inputs than ours. Del Boca et al. (2014) is the closest to us in terms of inputs and directly assumes complementarities by positing a Cobb-Douglas production function.

the qualitative patterns of child development by mother’s work status and education. As for child care use, the model matches the very low use by households with non-working mothers, but over predicts the use of high-quality centers on the part of working mothers of all educations. The price paid for child care services is well matched, both in terms of magnitude and ranking across mother’s education and work status. Relative to other moments, the fit for some moments of father’s wages is less tight. This was expected given that we identify wage offers by education for every parent only through mother’s education level.³¹

6 Policy Simulations

6.1 Set-up

We focus on interventions relevant to the current debate on how to increase access to center-based care (e.g., Chzhen et al. (2019)). The main tools to expand access are child care subsidies and public provision. In addition, recent evidence about the potentially harmful effects of low-quality care has led policymakers to focus on center quality, leading to the establishment of minimum quality standards.³²

Thus, we simulate four types of policies:

1. *vouchers*, which transfer cash that can only be used to buy child care services;
2. *cash transfers*, which transfer cash that can be used to buy child care services but also other goods or services;
3. *quality regulation*, which imposes a minimum level of quality on child care centers;
4. *public provision*, whereby only the public sector operates child care centers.

Using our parameter estimates, we first compute the baseline equilibrium (or “baseline”). Then we compute the equilibrium for each of these policies, and compare their outcomes with those from the baseline to estimate policy impacts. We focus on impacts on child development, child care use, child care supply, mothers’ labor market, and fiscal costs.

Because we use our equilibrium model to evaluate the policies, the estimated effects are *equilibrium effects*. To see their importance, consider, for instance, a policy that provides child care vouchers to all families. This policy would raise the demand for child care and, if the supply of child care centers did not change, it would raise the use and price of child care – as predicted by a partial equilibrium model. Yet our general equilibrium model allows for an endogenous adjustment of the supply – in this case, a supply increase as new providers enter the market encouraged by the demand expansion. As a result, we expect a greater expansion of child care use, and a lower price increase, than in partial equilibrium. These

³¹Recall from Section 4.1.1 that we fit moments conditioning only on the education of the mother.

³²About the effects of center’s quality and the policy reactions, see Baker et al. (2008) and Berlinski and Schady (2015). For a rare study on minimum quality standards, see Xiao (2010).

effects are particularly relevant in our context because the policies under consideration have been either implemented or proposed for a large share of households.³³ Therefore, they are more likely to trigger general equilibrium effects than policies concerned with small changes relative to the baseline.

We implement our policy simulations as follows. The *voucher* consists of a cash amount of 40 dollars a day per child, or five dollars per hour of service. This amount roughly corresponds to families’ expenditure in high-quality child care centers in the baseline model. If child care price is up to five dollars, the family does not pay anything out of pocket; otherwise it must cover the difference between the actual price and five dollars.³⁴ We simulate two voucher types: universal (available to all households) and targeted (available only to households that meet a given requirement). We impose two alternative requirements for targeted vouchers: a work requirement, whereby both parents must work to receive the voucher (i.e., the family must be a “working family”), and an income requirement, which limits the voucher to households with an income in the bottom quartile of the baseline household income distribution.³⁵ Parents take into account whether the voucher is universal or targeted when making their child care and labor supply decisions.

We set the *cash transfer* to the same amount of the voucher, namely 40 dollars a day per child. The transfer is universal and does not impose any requirement. While the household can use it to purchase child care services, it can also use it to consume more. This experiment mimics actual cash transfer programs to families with newborn and young children, and provides a useful comparison to vouchers.³⁶

In the *quality regulation* simulation, centers can only offer high-quality services if they wish to operate in the market (i.e., $\gamma_k = \gamma^H$ for all k). This experiment mimics standards-based policies that only allow high-quality centers to operate in the market.³⁷ Hence, at the entry stage of the model firms decide whether to enter as high-quality providers, or to not enter at all. A common concern with quality regulation is that high-quality centers might be costly, pricing out of the market the families that cannot afford them. In the *quality regulation plus voucher* simulation we implement a policy combination that might alleviate this concern: whereas only high-quality centers are allowed to operate, a (work requirement) voucher of five dollars per hour of service is offered to working families, in which both parents work.

Finally, the *public provision* simulation mimics a public entity, such as a local or national government, who is the sole provider of child care services. Since our model does not include a government, we simulate the existence of a benevolent policymaker who controls the supply of child care. We assume that this “government” faces a series of marginal costs for providing low and high quality - due, for instance, to the heterogeneous skills of centers’ managers and care providers. These cost pairs create the same heterogeneity as the firm types in our baseline

³³For example, the expansions of center-based care studied by Baker et al. (2008) and Havnes and Mogstad (2011a) apply to all children in the relevant age range of, respectively, Quebec and Norway.

³⁴If child care price is lower than five dollars, then the unused portion of the voucher returns to the policymaker. In equilibrium, however, prices are not lower than five dollars when vouchers are available.

³⁵The 25th percentile threshold corresponds to a daily income of 92 dollars.

³⁶Universal child benefits, frequently taking the form of actual cash transfers, are common in numerous European countries and Canada. See for example, OECD (2011), González (2013), and Schirle (2015).

³⁷For actual policy examples and evaluations, see Blau and Currie (2006), Hotz and Xiao (2011), and Xiao (2010).

model. We can interpret them as heterogenous individuals hired by the government in order to run the child care centers. The government considers all possible market configurations and chooses the one that maximizes social surplus, subject to the constraint of zero profits for each provider. In other words, each provider charges a price equal to its average cost. The taste for variety on the part of households implies that a given quality might be supplied by different providers, at different costs and hence prices. Importantly, we assume that families -and not the government- pay for the service. Indeed, even in countries with widespread public provision (e.g., Denmark, Germany, and Sweden), most families still pay. The public aspect of provision refers to the supply-side control exerted by the government, but does not entail free services. Rather, the supply-side control seeks to lower prices so that families can gain access to the service.

Formally, let $Z = \{(c_1^L, c_1^H), \dots, (c_K^L, c_K^H)\} \times \{0, \gamma^L, \gamma^H\}$ be the space of market configurations and $W(z, \mathbf{p})$ be the social surplus associated with the market configuration $z \in Z$ and the price vector \mathbf{p} . In equilibrium, the set of firms that provide child care, along with their qualities and prices, are those that maximize social surplus, or

$$\max_z W(z, \mathbf{p}) \quad \text{s.t.} \quad p_k = c_k + \frac{F}{N_k(\mathbf{p})}, \quad \forall k = 1, \dots, K, \quad (25)$$

where $N_k(\mathbf{p})$ denotes the total enrollment of provider k associated with the price vector \mathbf{p} .

To provide intuition for public provision, consider two extreme thought experiments. In the first, consider households that do not have idiosyncratic preferences or taste for variety. In this case, in principle three possible market configurations are possible depending on model parameters: a monopoly with the most efficient high-quality provider, a monopoly with the most efficient low-quality provider, and a duopoly with the most efficient high- and low-quality providers. These configurations have one element in common— they minimize the number of providers and hence spread their fixed costs among a large number of households, therefore reaching very low prices. In this thought experiment there is minimum variety, but minimum costs as well. In the second thought experiment, consider households that do have idiosyncratic preferences, but firms that do not have a fixed cost. In equilibrium there are many firms in this market - potentially one for each household. In this thought experiment there is maximum variety, but not minimum costs.

By maximizing social surplus subject to a zero-profit constraint for each provider, public provision strikes a balance between costs and variety, and is reminiscent of the social optimum in monopolistic competition (Dixit and Stiglitz, 1977). In general, then, public provision is expected to yield fewer firms and hence lower prices than the baseline.

6.2 Results

Table 5 reports our counterfactual results. Column 1 shows the baseline equilibrium, and columns 2 to 8 report the equilibrium for each policy simulation. Panels a, b, c, d, and e refer to impacts on child development, child care use, child care markets, mothers' labor market, and fiscal costs respectively. We use the terms *subsidy policies* to refer to vouchers (including quality regulation plus vouchers) and cash transfers, and *non-subsidy policies* for quality regulation (without vouchers) and public provision. Subsidy (except for quality

regulations plus vouchers) and non-subsidy policies seek to affect child care demand and supply, respectively. Quality regulation plus vouchers, in turn, seeks to affect both demand and supply.

Note that, while a family might be eligible for a subsidy, the family might choose not to take it. Hence, in panel b we specifically report the proportion of children who actually use the subsidy. In the discussion that follows, all comparisons are relative to the baseline, unless indicated otherwise. For expositional purposes, the discussion focuses on a specific set of impacts at a time. However, the fact that these are equilibrium effects implies that, in order to explain a given impact, we often need to refer to others as well.

6.2.1 Child development

As Table 5's panel (a) shows, all simulated policies yield, on average, child development gains. Most of them deliver sizable gains - larger than one standard deviation - for children of both high school and college mothers. Not surprisingly, the greatest average gains come from quality regulation plus vouchers, which at the same time raise center quality and subsidize high-quality centers for working families. Universal vouchers and cash transfers also deliver large quality gains, yet in different ways. Universal vouchers raise child development by raising child care use (panel b), whereas cash transfers raise child development by allowing more mothers to stay at home with their children instead of working (panel d).

Work requirement and income requirement vouchers deliver the smallest child development gains - lower than one standard deviation. Since work requirement vouchers provide an incentive for work, mothers work more (panel d) and spend less time with their children, thus leading to small child development gains. Income requirement vouchers, by design, could potentially benefit a quarter of the families. Given that only 14 percent of the families take up the voucher (panel b), they lead to small child development gains.

6.2.2 Child care use

Table 5's panel (b) shows that all policies, with the exception of income requirement vouchers, raise child care use. Public provision and universal vouchers raise it the most because they provide child care services at a low out-of-pocket cost to parents. In contrast, income requirement vouchers and quality regulation raise child care use the least - the former conditions access to vouchers; the latter provides no subsidy for child care and (slightly) raises its average price (panel c).

Although both the universal voucher and the cash transfer provide a five-dollar subsidy to all households, their effects on child care use are quite different - and different, in turn, for college and high school mothers. Child care use rises substantively under universal vouchers, both for high school and college mothers, because the voucher enables both sets of mothers to pay for child care and work more (panel d). In contrast, under cash transfers child care use rises for high school mothers (albeit less than with universal vouchers) yet falls for college mothers. The reason is that the cash transfer enables all mothers to work less (panel d) and spend more time with their children. College educated mothers, who are more productive than even high-quality centers in the production of child development, lower their child care

use. The opposite happens to high school mothers, who are less productive than high-quality centers yet can now afford them through the cash transfer.

6.2.3 Child care markets

Table 5's panel (c) reports the effects of policies on child care markets. The first row shows the proportion of active providers relative to all potential entrants and is thus a measure of the number of providers in the market. This number rises or stays the same with subsidy policies, which specifically subsidize demand. However, it drops under non-subsidy policies, which either allow only high-quality centers to operate, or that concentrate supply among few providers. Although non-subsidy policies lower the number of firms in the market, they attain the highest fraction of high-quality centers because they concentrate supply among them.³⁸

Changes in child care use as well as in the number of centers lead to changes in average enrollment per center. This rises the most under public provision, which features high child care use and a low number of firms. It drops the most for cash transfers, which features a large number of firms and low child care use. In general, though, our simulated policies lead to an increase in average enrollment, suggesting that at least some of the additional demand is absorbed by the existing centers.

Changes in child care demand and supply lead to changes in average price as well. With the exception of public provision, which provides child care at cost, average price rises in all simulated policies, both for low and high-quality centers. This price increase is the net effect of three forces. The first is higher demand, which alone would lead to higher prices. The second is higher supply (except in the case of quality regulation), which alone would lead to lower prices. The third is supply composition. Since new entrants have higher marginal costs than incumbent providers, they charge higher prices. The net price increase is thus driven by higher demand and by the entry of less efficient providers.

Policymakers often fear that, while quality regulation can indeed raise quality, it can also raise prices as high-quality services are more costly. Our simulations do show an average price increase under quality regulation, yet a very slight one. Further, this price increase is driven by a composition effect due to the elimination of low-quality firms from the market. When we focus on the average price for high-quality firms, we see that it actually falls. The reason is that, by allowing only high-quality centers to operate, this policy effectively increases the number of high-quality firms in the market, which must then compete in prices. This price decline, however, is not enough to encourage greater child care use than in the baseline (panel b). Policymakers' fears that some baseline users of child care would lose access are thus not realized, but their hopes that additional users would gain access are not realized either.

In contrast, quality regulation plus vouchers does expand child care access (panel b) while maintaining high quality standards. Average prices rise relative to the baseline because quality is higher and demand rises due to the voucher. However, supply rises as well, and family incomes rise as more mothers work. As a result, child care use rises, both for high

³⁸The predominance of high-quality providers under public provision is not a general result. In our counterfactuals, it is largely driven by the fact that the marginal cost difference between high- and low-quality providers is relatively small.

school and college mothers.

6.2.4 Mothers' labor market

Table 5's panel (d) reports mothers' labor market. Unlike child care use, which rises for all but one policy, mothers' labor force participation rises only in about half of the policies and falls in the others. It rises the most for work requirement vouchers and for quality regulations plus vouchers, given that both policies provide a clear incentive to work while also subsidizing child care. Quality regulation alone barely affects labor force participation given its slight effect on prices, but when combined with the work requirement voucher it leads to more working mothers. The unconditional child care subsidy embedded in universal vouchers also leads to a substantial increase in the number of working mothers. While universal and work requirement vouchers expand the labor force participation of college mothers by twice as much as for high school mothers, the effects on all mothers are quite similar under quality regulation plus voucher. The reason is that the latter policy improves child care quality proportionally more for high school mothers, who are thus more compelled to work in order to qualify for the voucher.

In contrast, labor force participation falls under cash transfers, income requirement vouchers, and public provision. Cash transfers leads many mothers to work less and care for their children at home. The effect is greater for college mothers than for high school mothers given their greater productivity in child development production. Income requirement vouchers give families an incentive to keep their income below the voucher eligibility threshold. This discourages mothers' work, whose additional earnings might entail the loss of voucher eligibility. Once again, the decline is greater for college mothers, who are more productive in child development production than high school mothers. Among the mothers that "switch" to staying at home, high school mothers use child care more, whereas college mothers use it less (panel b).

Finally, public provision lowers labor force participation because it provides high-quality child care at such low prices that many mothers no longer need to work. The decline in labor force participation is greater among high school than college mothers, since the former are less productive in child development production and hence benefit more from high-quality child care. Interestingly, although public provision raises child care use the most and leads to a very high fraction (almost 90 percent) of high-quality centers, it does not deliver the highest child development gains. The reason is that the availability of low cost, high-quality child care crowds out mothers' time with their children. For the children of college educated mothers, this entails a decline in child development.

With the exception of cash transfers and quality regulation, all policies raise mothers' average wages and thus shrink the gender wage gap. In particular, universal vouchers completely eliminate the gender wage gap since the additional mothers who join the labor force have higher reservation wages than those participating at the baseline. A similar mechanism, yet to a lesser degree, operates for the other policies that shrink the gender wage gap.

6.2.5 Fiscal cost

Ultimately the policymaker might need to choose a policy among those presented here. This choice would depend, in part, on the fiscal cost of the policies. We provide some metrics of fiscal costs and report them in panel (e) of Table 5. In the case of *subsidy policies*, it is relatively straightforward to calculate fiscal costs - they consist of the total dollar amount of subsidies distributed to families, abstracting from any administrative costs. Fiscal costs are harder to calculate for *non-subsidy policies*. Under quality regulation, administrative costs (for instance, of high quality enforcement) are the main fiscal cost, but we lack data on them. Public provision also entails administrative costs (for instance, of identifying the relevant providers), but we lack data on these as well. Hence, we do not report fiscal costs for non-subsidy policies.

Given our lack of data on administrative costs, our fiscal cost calculations are a first-order approximation to the “true” fiscal costs. Moreover, they do not take into account the fiscal benefits or losses due to changes in mothers’ labor force participation, or the future fiscal savings derived from early gains in child development (see, Currie (2001) and Heckman et al. (2010b)).

With these caveats in mind, we report three metrics for fiscal costs: per-hour monetary cost per child in the economy (regardless of whether he/she uses the policy), total cost relative to families’ total income, and total cost relative to child care providers’ total profits. We refer to the last two as “implicit tax rates” because they approximate the tax rates that families and child care centers, respectively, would face in order to finance the policy.

By far, the most costly policy is cash transfers, as it provides a subsidy to every family in the economy (panel b). The fiscal cost of the cash transfer amounts to almost a quarter of centers’ total profits and 18 percent of families’ total labor income. Universal vouchers are costly but only half as much, since they are provided only to families who use outside child care. The least costly policy is income-requirement vouchers, which benefit the smallest number of families (panel b).

Which policy is preferable depends on the policymakers’ objective. If the objective is to maximize child care use, then public provision is the best policy. Indeed, by its very design public provision is the policy that maximizes social surplus. Relative to the baseline, child care use rises and so does child development; prices fall, and the number of high-quality centers rises as well.

Nonetheless, public provision leads to a decline in mothers’ labor force participation. If the objective is to maximize mothers’ labor force participation at the lowest possible fiscal cost, then work requirement vouchers, and quality regulation plus vouchers, appear as the best policies. The latter, however, delivers much higher child development gains – the highest, in fact, of all the policies under consideration.

It appears, then, that a policymaker aiming at expanding mothers’ labor force participation while ensuring high-quality child care at a low fiscal cost should gravitate towards quality regulation plus vouchers. This policy succeeds by *combining* demand-side and supply-side instruments. Unlike public provision, this policy relies on market forces to provide high-quality care to all children in child care centers. It provides more options to families than public provision, as it features more providers, and thus satisfies families’ taste for variety. Through

the voucher, it gives families the ability to afford high-quality care. If we were to take into account the fiscal benefits derived from the increase in mothers' labor force participation due to the voucher component, these conclusions would only be stronger.

6.3 Comparison with Previous Literature

Universal vouchers' simulated effects can be compared with those described in Baker et al. (2008), who study the impact of a major child care subsidy in the Canadian province of Quebec.³⁹ As in our simulations, they find an increase in child care use and female labor supply. Contrary to our simulations, they find no effect on children's cognitive outcomes, presumably because the subsidy was used by relatively affluent children who would otherwise have stayed home with their mothers. These children are likely to have highly educated mothers, who could have provided child care of a quality equal or superior to the one provided by the centers. A similar result is also implied by Fort et al. (2019), who study a program that led to a preschool enrollment increase for 2 year-old children in Bologna, Italy. Our theoretical model contains all the channels that could have generated results similar to those in Baker et al. (2008) and Fort et al. (2019). However, our simulation results are different likely due to two elements: we allow for an endogenous adjustment of supply in response to the increased demand due to the policy; and a broad set of families use the vouchers, including some high school and college mothers for whom their chosen child care center represents a child development improvement.

We can compare the effects of our targeted voucher simulations with those of several actual programs. For example, Bettendorf et al. (2015) study a child care policy in the Netherlands, in which the subsidy amount is decreasing in family income and covers between 25 and 95 percent of child care costs. They find only modest increases in maternal employment rate, a result consistent with the relatively small decline in mothers' participation rates implied by our income-targeted voucher simulations.

Some tax credit policies are similar to our voucher simulations with income and work requirements. The fungibility of the credit is different, since tax credits can be used to buy any goods and services, but the conditionality is the same. For instance, the Earned Income Tax Credit (EITC) requires that at least one household member work, and the benefit amount received decreases after a certain income threshold. There is evidence that EITC raises labor market participation for single mothers but lowers it for secondary earners (Eissa and Hoynes (2006) and Hotz (2003)). Dahl and Lochner (2012), Dahl and Lochner (2017) and Lundstrom (2017) find that EITC expansions improve children's math and reading achievement. However, Heckman and Mosso (2014) caution against taking these results as evidence that cash transfers significantly improves children's outcomes because they ignore equilibrium effects on female labor supply. Agostinelli and Sorrenti (2018) improve on this margin by controlling for the endogeneity of family income and maternal labor supply to study the EITC and its impact on single mothers who work. They find that a subsidy of \$1,000 a year increases average child development by about 4.4 percent of a standard deviation. The magnitude is comparable to our estimated effects on families with high school mothers, since

³⁹All parents were offered subsidized childcare at the cost of \$5 per day for children aged 0 to 4. Supply was adjusted by expanding center-based care, as in our experiments.

our simulated subsidy is about ten times larger, and our effect on average child development is slightly more than ten times larger.⁴⁰ Our slightly larger impact is due to two main differences relative to their study: we focus on two-parent households, and we allow both labor supply and child care supply to adjust in response to the policy.

Since the children in our sample are very young, our cash transfer simulation may be compared to cash transfers for families with newborn or very young children, which are common in numerous European countries and Canada (OECD, 2011). González (2013) considers the effect of a significant and unanticipated lump-sum universal child cash transfer to mothers with newborn babies. She finds that the policy led to a significantly lower use of formal day care services and a temporary reduction in mother’s labor supply, leading to an increase in maternal care time. Schirle (2015) studies the Canadian Universal Child Care Benefit, a program distributing a monthly amount (\$100 in 2012) for each child under six year of age. She finds substantial reductions in labor supply both for mothers with high and low education levels, although the impacts are stronger for the latter. Milligan and Stabile (2009) studies the expansion of transfers to families with children ages zero to five in the Canadian state of Manitoba. They find large differences in labor supply responses by mother’s education, and estimate a negative impact only for families with low education levels. In our cash transfer simulation, we also find different effects on labor force participation and child development by mother’s education. Effects on child development (as measured by our child development measure) are larger for college mothers due to their greater productivity in child development production. Since education is correlated with income, these heterogenous effects by mother’s education are consistent with the heterogenous effects by income estimated in the literature, as exemplified by Løken et al. (2012) using Norwegian data.

Our quality regulation simulation connects with an established literature that has identified key trade-offs.⁴¹ Blau (2003) and Hotz and Xiao (2011) articulate the most important one: while quality regulations improve quality, they frequently lower use because higher quality is more expensive. The negative impact of lower use is compounded by the selection of the children priced out of the market -children who come from families with low income and low parental education, and who would benefit the most from high-quality child care. Hotz and Xiao (2011) find that, in the US, state-level regulations on the supply and quality of child care services has led to fewer centers but higher quality for the remaining ones. Consistent with the aforementioned trade-off, the reduction in the number of centers is concentrated in low-income areas, while the quality increase is concentrated in high-income areas.

In our quality regulation simulation, this trade-off almost disappears, as enrollment remains stable for all families due to the modest average price increase. On average, families that pay \$3.02 per hour on low-quality centers in the baseline pay more (\$4.69) for high-quality centers in the counterfactual. However, this is less than what they would pay for a high-quality center in the baseline (\$5.59). On average, families that pay \$5.59 per hour on high-quality centers in the baseline pay less (\$4.69) for those centers in the counterfac-

⁴⁰We compare with families with high school mothers since they are most similar to the sample used in Agostinelli and Sorrenti (2018). In our simulations, the subsidy is between \$10,000 and \$14,000 a year depending on how many days the service is used, and the average effects for the income requirement voucher is about 55 percent of a standard deviation.

⁴¹For a survey, see Blau and Currie (2006).

tual. The reason for the relatively low price of high-quality centers in the counterfactuals is that, due to the regulation, new high-quality providers enter the market and push prices downward. This is an additional example of how our counterfactuals, which incorporate supply-side equilibrium responses, lead to different conclusions than those from a partial equilibrium analysis. The simulation combining quality regulation with a work requirement voucher only strengthens this point, as the price of high-quality centers barely increases but enrollment rises substantially, leading to large gains in child development.

An interesting result of our public provision simulation is that, although it maximizes child care use and raises the proportion of high-quality child care centers substantially, it does not raise average child development as much as the cash transfer. The reason is similar to that argued by Baker et al. (2008) and Fort et al. (2019) - some of the mothers using public provision could, under the cash transfer, provide care of even higher quality than that provided by high-quality centers.

7 Conclusions

Child care markets are at the center of the trade-off facing parents as they allocate time and material resources. Material and time inputs are both essential to successful child development. By supplying labor in the market, parents obtain the income necessary to provide material inputs, yet also lose a portion of the time that would otherwise be available to care for their children. Therefore, access to high-quality centers, which can effectively substitute for parents' time, is crucial to the debate on family-friendly policies. Any policy seeking to expand access to center-based care should take into account that center quality and price are market equilibrium outcomes, which will adjust in response to policy interventions.

We model the child care market and estimate its structural parameters, endogenizing both demand and supply in order to evaluate policies in an equilibrium context. We model parents' choices about labor supply, child care, and time allocation, as well as providers' decisions on entry, quality, and price. In equilibrium, the market features multiple levels of quality, price dispersion, and horizontal differentiation. Data on child care markets, child care use, child cognitive development, and parental labor supply allows us to identify the model parameters. We estimate the model using a Simulated Method of Moments procedure on a sample of two-year old children from the 2003 Early Childhood Longitudinal Study (ECLS-B), a nationally representative sample of children followed from birth through kindergarten. We obtain precise point estimates that provide a good fit of the data and are comparable to others in the literature.

We use the estimated model to simulate child care policies that provide a demand subsidy (universal vouchers, vouchers with an income or work requirement, and cash transfers) or that affect supply (regulation of center quality, and public provision). Our counterfactual results illustrate the importance of taking equilibrium effects into account in policy evaluation. For example, policymakers often worry that, while quality regulations can indeed raise child care quality, they may also raise prices, thus depriving the neediest children of the service.⁴² While our simulations do report a price increase as a result of imposing quality regulation, they

⁴²For studies on this issue, see Blau and Currie (2006), Blau (2003) and Hotz and Xiao (2011).

show it to be small. The reason is that, by endogenizing the supply of high-quality providers, the policy stimulates their entry and leads to lower prices that those that would otherwise prevail. As a result, child care use remains almost the same, but its quality rises. Equally interesting is the result obtained from a policy combining quality regulation with vouchers conditional on work requirements. Of all the policies considered, this one leads to the greatest gains in child development while at the same time expanding child care use and female labor supply at a low fiscal cost.

Ultimately, which policy is preferable depends on the policymaker's objective - whether to maximize child care use, child development, or female labor force participation - as well as fiscal resources. The range of our simulation results illustrates the tension among different policy objectives, and empirically shows how the endogenous adjustments in the child care market can magnify or dampen the policies' adverse effects. It also illustrates the effectiveness of combining demand- and supply-side instruments, as in the case of quality regulation plus work requirement vouchers, rather than seeking to affect only one side of the market. Indeed, the policymaker can mitigate the unintended effects of one instrument by combining it with another. An equilibrium analysis, such as the one presented in this paper, is indispensable for the identification of such unintended effects and for the design of sound policy.

References

- Agostinelli, F. and Sorrenti, G. (2018). Money vs. time: Family income, maternal labor supply, and child development. Technical Report 273.
- Agostinelli, F. and Wiswall, M. (2016). Estimating the technology of children’s skill formation. Technical report, National Bureau of Economic Research.
- Attanasio, O., Cattan, S., Fitzsimons, E., Meghir, C., and Rubio-Codina, M. (2015a). Estimating the production function for human capital: Results from a randomized control trial in colombia. In *Working Paper 20965*. NBER.
- Attanasio, O., Meghir, C., and Nix, E. (2015b). Human capital development and parental investment in india. Technical report, National Bureau of Economic Research.
- Baker, M., Gruber, J., and Milligan, K. (2008). Universal child care, maternal labor supply, and family wellbeing. *Journal of Political Economy*, 116(4):709–745.
- Bastos, P. and Cristia, J. (2012). Supply and quality choices in private child care markets: Evidence from são paulo. *Journal of Development Economics*, 98(2):242–255.
- Berlinski, S. and Schady, N. (2015). *The early years: child well-being and the role of public policy*. Springer.
- Berlinski, S. and Vera-Hernández, M. (2019). The economics of early interventions aimed at child development. In *Oxford Research Encyclopedia of Economics and Finance*.
- Bernal, R. (2008). The effect of maternal employment and child care on children’s cognitive development. *International Economic Review*, 49(4):1173–1209.
- Bernal, R. and Keane, M. P. (2011). Child care choices and childrens cognitive achievement: The case of single mothers. *Journal of Labor Economics*, 29(3):459–512.
- Bettendorf, L. J., Jongen, E. L., and Muller, P. (2015). Childcare subsidies and labour supply: Evidence from a large dutch reform. *Labour Economics*, 36:112–123.
- Björklund, A. and Salvanes, K. G. (2011). Education and family background: Mechanisms and policies. In *Handbook of the Economics of Education*, volume 3, pages 201–247. Elsevier.
- Black, S. E. and Devereux, P. J. (2010). Recent developments in intergenerational mobility. Technical report, National Bureau of Economic Research.
- Blau, D. and Currie, J. (2006). Pre-school, day care, and after-school care: Who’s minding the kids? *Handbook of the Economics of Education*, 2:1163–1278.
- Blau, D. M. (2003). Do child care regulations affect the child care and labor markets? *Journal of Policy Analysis and Management*, 22(3):443–465.
- Blau, D. M. and Mocan, H. N. (2002). The supply of quality in child care centers. *Review of Economics and Statistics*, 84(3):483–496.
- Blau, D. M. and Robins, P. K. (1988). Child-care costs and family labor supply. *The Review of Economics and Statistics*, pages 374–381.
- Campbell, F., Conti, G., Heckman, J. J., Moon, S. H., Pinto, R., Pungello, E., and Pan, Y. (2014). Early childhood investments substantially boost adult health. *Science*,

- 343(6178):1478–1485.
- Carneiro, P., Lken, K. V., and Salvanes, K. G. (2015). A flying start? Maternity leave benefits and long-run outcomes of children. *Journal of Political Economy*, 123(2):365–412.
- Carneiro, P., Meghir, C., and Parey, M. (2013). Maternal education, home environments, and the development of children and adolescents. *Journal of the European Economic Association*, 11(suppl_1):123–160.
- Chzhen, Y., Rees, G., and Gromada, A. (2019). Are the world’s richest countries family friendly? Policy in the OECD and EU. *UNICEF Office of Research*.
- Ciliberto, F. and Tamer, E. (2009). Market structure and multiple equilibria in airline markets. *Econometrica*, 77(6):1791–1828.
- Conti, G., Heckman, J. J., and Pinto, R. (2016). The effects of two influential early childhood interventions on health and healthy behaviour. *The Economic Journal*, 126(596):F28–F65.
- Cunha, F., Heckman, J. J., and Schennach, S. M. (2010). Estimating the technology of cognitive and noncognitive skill formation. *Econometrica*, 78(3):883–931.
- Currie, J. (2001). Early childhood education programs. *Journal of Economic perspectives*, 15(2):213–238.
- Currie, J. and Almond, D. (2011). Human capital development before age five. In *Handbook of labor economics*, volume 4, pages 1315–1486. Elsevier.
- Dahl, G. B. and Lochner, L. (2012). The impact of family income on child achievement: Evidence from the earned income tax credit. *American Economic Review*, 102(5):1927–56.
- Dahl, G. B. and Lochner, L. (2017). The impact of family income on child achievement: Evidence from the earned income tax credit: Reply. *American Economic Review*, 107(2):629–31.
- Del Boca, D., Flinn, C., and Wiswall, M. (2014). Household choices and child development. *The Review of Economic Studies*, 81(1):137–185.
- Dixit, A. K. and Stiglitz, J. E. (1977). Monopolistic competition and optimum product diversity. *The American economic review*, 67(3):297–308.
- Dranove, D., Gron, A., and Mazzeo, M. J. (2003). Differentiation and competition in hmo markets. *The Journal of Industrial Economics*, 51(4):433–454.
- Dustmann, C. and Schonberg, U. (2012). Expansions in maternity leave coverage and children’s long-term outcomes. *American Economic Journal: Applied Economics*, 4(3):190–224.
- Eissa, N. and Hoynes, H. W. (2006). Behavioral responses to taxes: Lessons from the eitc and labor supply. *Tax policy and the economy*, 20:73–110.
- Flinn, C. and Heckman, J. (1982). New methods for analyzing structural models of labor force dynamics. *Journal of Econometrics*, 18(1):115 – 168.
- Fort, M., Ichino, A., and Zanella, G. (2019). Cognitive and non-cognitive costs of daycare 0-2 for girls. *Journal of Political Economy*.
- González, L. (2013). The effect of a universal child benefit on conceptions, abortions, and early maternal labor supply. *American Economic Journal: Economic Policy*, 5(3):160–88.

- Grimshaw, D. and Rubery, J. (2015). The motherhood pay gap: a review of the issues, theory and international evidence. *ILO Conditions of Work and Employment Series No. 57*.
- Hart, O. (1995). *Firms, contracts, and financial structure*. Clarendon press.
- Havnes, T. and Mogstad, M. (2011a). Money for nothing? universal child care and maternal employment. *Journal of Public Economics*, 95(11-12):1455–1465.
- Havnes, T. and Mogstad, M. (2011b). No child left behind: Subsidized child care and children’s long-run outcomes. *American Economic Journal: Economic Policy*, 3(2):97–129.
- Heckman, J., Moon, S. H., Pinto, R., Savelyev, P., and Yavitz, A. (2010a). Analyzing social experiments as implemented: A reexamination of the evidence from the highscope perry preschool program. *Quantitative economics*, 1(1):1–46.
- Heckman, J. J. (1974). Effects of child-care programs on women’s work effort. In *Marriage, Family, Human Capital, and Fertility*, pages 136–169. NBER.
- Heckman, J. J., Krueger, A. B., et al. (2005). Inequality in america: What role for human capital policies? *MIT Press Books*, 1.
- Heckman, J. J., Moon, S. H., Pinto, R., Savelyev, P. A., and Yavitz, A. (2010b). The rate of return to the highscope perry preschool program. *Journal of public Economics*, 94(1-2):114–128.
- Heckman, J. J. and Mosso, S. (2014). The economics of human development and social mobility. *Annu. Rev. Econ.*, 6(1):689–733.
- Helburn, S. W. and Howes, C. (1996). Child care cost and quality. *The future of children*, pages 62–82.
- Hotz, V. J. (2003). The earned income tax credit. In *Means-tested transfer programs in the United States*, pages 141–198. University of Chicago press.
- Hotz, V. J. and Xiao, M. (2011). The impact of regulations on the supply and quality of care in child care markets. *The American Economic Review*, pages 1775–1805.
- Kimmel, J. and Connelly, R. (2007). Mothers’ time choices caregiving, leisure, home production, and paid work. *Journal of Human Resources*, 42(3):643–681.
- Laughlin, L. (2013). *Who’s minding the kids?: Child care arrangements, Spring 2011*. US Department of Commerce, Bureau of the Census.
- Løken, K. V., Mogstad, M., and Wiswall, M. (2012). What linear estimators miss: The effects of family income on child outcomes. *American Economic Journal: Applied Economics*, 4(2):1–35.
- Lundstrom, S. (2017). The impact of family income on child achievement: Evidence from the earned income tax credit: Comment. *American Economic Review*, 107(2):623–28.
- Mazzeo, M. J. (2002). Product choice and oligopoly market structure. *RAND Journal of Economics*, 33(2):1–22.
- Milligan, K. and Stabile, M. (2009). Child benefits, maternal employment, and children’s health: Evidence from canadian child benefit expansions. *American Economic Review*, 99(2):128–32.

- Mitchell, A. and Stoney, L. (2012). *Generic Cost Model: Centers*. Alliance for Early Childhood Finance.
- OECD (2011). *Doing Better for Families*.
- Reiss, P. C. and Wolak, F. A. (2007). Structural econometric modeling: Rationales and examples from industrial organization. *Handbook of econometrics*, 6:4277–4415.
- Ribar, D. C. (1992). Child care and the labor supply of married women: Reduced form evidence. *Journal of human resources*, pages 134–165.
- Ruhm, C. J. (2004). Parental employment and child cognitive development. *The Journal of Human Resources*, 39(1):155–192.
- Schirle, T. (2015). The effect of universal child benefits on labour supply. *Canadian Journal of Economics/Revue canadienne d'économique*, 48(2):437–463.
- Schultz, T. P. (1993). Returns to women's education. *Women's education in developing countries: Barriers, benefits, and policies*, pages 51–99.
- Seim, K. (2006). An empirical model of firm entry with endogenous product-type choices. *The RAND Journal of Economics*, 37(3):619–640.
- UNICEF (2017). Building the future: Children and sustainable development goals in rich countries. *UNICEF Innocenti Report Card 14*.
- Urquiola, M. and Verhoogen, E. (2009). Class-size caps, sorting, and the regression-discontinuity design. *The American Economic Review*, 99(1):179–215.
- Xiao, M. (2010). Is quality accreditation effective? evidence from the childcare market. *International Journal of Industrial Organization*, 28(6):708–721.

Table 1: Descriptive Statistics

Variables	Mean	Std. Dev.
<u>Child Care Market:</u>		
Use: Child in Center	30.90	44.70
Definition: =1 if main form of care is a day care center		
Units: Percentage		
Price: Center Fee	4.63	1.58
Definition: Hourly fee for a 2 year old attending full-time		
Units: 2003 U.S. dollars per hour		
Quality: Child in High Quality Center	4.17	0.92
Definition: Average overall ITERS score		
Units: ITERS scale 1–7		
<u>Child’s Cognitive Score:</u>		
Definition: Bailey Short Form Mental Ability Score	12.88	1.13
Units: Score/10		
<u>Parents’ Labor Market:</u>		
<u>Mothers</u>		
Proportion Working	27.50	19.90
Definition: =1 if mother works full-time for pay		
Units: Percentage		
Wage	21.33	1.40
Definition: Gross hourly wage in 2003 dollars		
Units: 2003 U.S. dollars per hour		
Education	48.60	50.00
Definition: =1 if mother completed college or more		
Units: Percentage		
<u>Fathers</u>		
Proportion Working	100.00	0.00
Definition: =1 if father works full-time for pay		
Units: Percentage		
Wage	23.87	17.17
Definition: Gross hourly wage in 2003 dollars		
Units: 2003 U.S. dollars per hour		
Education	51.70	50.00
Definition: =1 if father completed college or more		
Units: Percentage		
<u>Sample Size</u>		
Number of Households	450	

NOTE: We report definitions of variables and descriptive statistics on the estimation sample extracted from the second wave of ECLS-B for children born in 2001. Sample size is rounded to the nearest 50 in order to comply with ECLS-B confidentiality rules.

Table 2: SMM Parameter Estimates

	Estimates	Standard Errors
<u>Panel a. Household Utility Function:</u>		
μ_{v1}	0.626	0.092
μ_{v2}	0.673	0.013
μ_{v3}	-0.950	0.061
σ_{v1}	0.242	0.099
σ_{v2}	0.395	0.320
σ_{v3}	0.168	0.053
μ_{ε}	0.816	0.153
<u>Panel b. Child Development Production Function:</u>		
γ_0	0.222	0.025
γ_f	0.052	0.004
γ_m	0.195	0.028
γ'_m	0.222	0.006
γ^L	0.179	0.028
γ^H	0.219	0.011
r	0.723	0.126
<u>Panel c. Parents Wage Offers:</u>		
μ_{w_f}	1.903	0.171
σ_{w_f}	0.480	0.021
μ'_{w_f}	3.390	0.451
σ'_{w_f}	0.532	0.156
μ_{w_m}	1.891	0.982
σ_{w_m}	0.414	0.043
μ'_{w_m}	2.959	0.014
σ'_{w_m}	0.548	0.025
ρ	0.630	0.180
ρ'	0.772	0.281
<u>Panel d. Centers Variable Costs:</u>		
μ_{c^L}	0.040	0.122
μ_{c^H}	0.423	0.028
σ_{c^L}	0.162	0.191
σ_{c^H}	0.201	0.089

NOTE: See main text for the parameters' definition. Bootstrap standard errors reported.

Table 3: Implied Distributions

	Expected Value	Standard Deviation
<u>Panel a. Household Utility Parameters:</u>		
$\alpha_c(j)$	0.360	0.073
$\alpha_q(j)$	0.374	0.093
$\alpha_f(j)$	0.075	0.018
$\alpha_m(j)$	0.191	0.033
<u>Panel b. Parents Wage Offers:</u>		
Overall:		
Fathers	22.74	11.28
Mothers	14.87	9.85
Both Parents HS:		
Fathers	17.92	7.07
Mothers	7.19	5.36
Both Parents C:		
Fathers	26.16	11.89
Mothers	21.73	12.48
Father C, Mother HS:		
Fathers	24.27	10.86
Mothers	11.62	8.14
Father HS, Mother C:		
Fathers	18.52	9.46
Mothers	21.43	12.60
<u>Panel c. Centers Variable Costs:</u>		
Low Quality	1.055	0.172
High Quality	1.558	0.316

NOTE: Definition: C is four-year college completed or more; HS is all other educational attainments. The distribution of the Utility Parameters is obtained by using the Household Utility Function parameters reported in Table 2 and applying Equations (16)–(19) in the main text. The distribution of the Wage Offers is obtained by using the Parents Wage Offers parameters reported in Table 2 and applying the definition of the lognormal distribution’s moments. The distribution of the Variable Costs is obtained by using the Centers Variable Costs parameters reported in Table 2 and applying the definition of the lognormal distribution’s moments.

Table 4: Model Fit: Sample and Simulated Moments

Moments	Mother	High School Mothers		College Mothers	
	Working	Sample	Simulated	Sample	Simulated
<u>Child Care Market:</u>					
Use: Child in Center	Y	1.00	1.00	1.00	1.00
	N	0.07	0.09	0.07	0.11
Price: Average	Y	3.93	4.03	4.78	4.81
	N	4.39	4.43	5.76	5.83
Price: Standard Deviation	Y	1.01	0.76	1.58	1.21
	N	2.05	1.54	1.98	1.74
Quality: Child in HQ Center	Y	0.53	0.63	0.46	0.55
	N	0.75	0.70	0.50	0.49
<u>Child Development:</u>					
Average	Y	12.84	12.71	13.36	13.22
	N	12.53	12.51	13.10	13.08
Standard Deviation	Y	1.07	1.05	0.87	0.92
	N	1.10	1.13	1.19	1.16
<u>Parents' Labor Market:</u>					
<u>Mothers</u>					
Proportion Working		0.18	0.22	0.38	0.41
Wage: Average		15.14	15.16	24.20	23.12
Wage: Standard Deviation		9.90	11.91	13.96	15.49
<u>Fathers</u>					
Proportion Working		1.00	1.00	1.00	1.00
Wage: Average	Y	17.66	17.85	26.36	26.75
	N	18.37	20.17	31.95	30.98
Wage: Standard Deviation	Y	8.74	10.07	14.33	16.10
	N	16.11	16.41	18.59	19.18

NOTE: *Sample* refers to the sample of 450 households used in estimation and extracted from the second wave of the ECLS-B for children born in 2001. *Simulated* refers to moments calculated from the equilibrium of our model, computed at the parameter estimates reported in Table 2. Y means Yes and N means No in reference to the presence of a working mother in the household. All the price and wage variables are in 2003 dollars per hour.

Table 5: Policy Experiments

	Baseline	Voucher – Conditionality:			Cash Transfer	Quality Regulation:		Public Provision
		None	Work	Income		Only	Work Voucher	
<u>Panel (a): Child development</u>								
Average child development gain								
All mothers	n.a.	1.47	0.45	0.43	1.83	1.42	2.21	1.08
High school mothers	n.a.	1.11	0.33	0.55	1.70	1.05	2.09	0.96
College mothers	n.a.	1.87	0.58	0.29	1.98	1.82	2.34	1.22
<u>Panel (b): Child care use</u>								
Proportions among:								
All mothers	0.38	0.56	0.52	0.37	0.40	0.39	0.54	0.57
High school mothers	0.29	0.45	0.44	0.34	0.35	0.29	0.45	0.44
College mothers	0.47	0.67	0.69	0.41	0.45	0.49	0.63	0.72
Proportion of children with subsidy:								
All mothers	0.00	0.56	0.45	0.14	1.00	0.00	0.44	0.00
High school mothers	0.00	0.45	0.31	0.13	1.00	0.00	0.36	0.00
College mothers	0.00	0.67	0.60	0.16	1.00	0.00	0.53	0.00
<u>Panel (c): Child care markets</u>								
Proportion of providers entering	73.33	80.00	80.00	73.33	86.67	60.00	73.33	53.33
Proportion of HQ centers	27.27	25.00	25.00	27.27	15.38	100.00	100.00	87.50
Average enrollment per center	28.77	38.77	36.15	28.18	25.46	35.82	40.63	59.71
Average price								
At all centers	4.55	6.57	6.04	5.02	5.26	4.69	5.86	2.53
At low-quality centers	3.02	5.12	4.37	3.13	3.21	n.a.	n.a.	1.84
At high-quality centers	5.59	7.23	6.93	6.10	5.71	4.69	5.86	2.68
<u>Panel (d): Mothers' labor market</u>								
Proportion mothers working								
All mothers	0.31	0.42	0.45	0.23	0.22	0.32	0.44	0.28
High school mothers	0.22	0.29	0.31	0.18	0.18	0.22	0.36	0.17
College mothers	0.41	0.56	0.60	0.29	0.26	0.42	0.53	0.39
Average wage								
All mothers	19.03	24.60	20.81	21.62	17.45	18.58	19.92	20.90
High school mothers	15.16	19.87	17.66	18.39	14.74	14.62	16.85	18.46
College mothers	23.12	29.60	24.15	25.04	20.31	22.78	23.17	23.48
<u>Panel (e): Fiscal cost</u>								
Per child in the economy	n.a.	2.79	2.26	0.71	5.00	n.a.	2.21	n.a.
Implicit tax rate on families	n.a.	0.08	0.07	0.02	0.18	n.a.	0.07	n.a.
Implicit tax rate on childcare centers	n.a.	0.08	0.07	0.06	0.24	n.a.	0.07	n.a.

NOTE: For a detailed description of the policy experiments see Section 6.1. All the price and wage variables are in 2003 U.S. dollars per hour. Definitions: *Baseline* is the model simulated at the estimated parameters; *HQ* denotes High Quality; *Average child development gain* measures the average difference in q between the counterfactual and the Baseline, standardized by the mean and standard deviation of the q distribution at Baseline; *Implicit tax rate on families* is the total cost of the policy relative to total income of all families; *Implicit tax rate on childcare centers* is the total cost of the policy relative to total profits of all centers. Finally, *n.a.* means not applicable or not available.

Appendices

A Simulation details

In this appendix we describe the quantitative version of our model. Our sample includes 450 households. We assume that each of them represents a household type, which yields $J = 450$ types. In our empirical application, a household type is a combination of parental wage offers, preferences, initial child development, and parental education. Hence, we construct the types as follows. First, for each type (household in the data) we take from the data the initial child development, father's education, and mother's education. Second, for each type, and for a given parameter vector, we generate parental wage offers and preference parameters. In particular, we draw the vector (v_1, v_2, v_3) from a joint normal distribution to create the vector of preference parameters, $(\alpha_c, \alpha_q, \alpha_f, \alpha_m)$, as defined in equations (16)-(19). Similarly, we draw a pair of wage offers, (w_f, w_m) from a joint lognormal distribution with the mean vector and variance-covariance matrix defined in (20) and (21), respectively.

We set the size of each household type and the non-labor income as follows: $M_j = 2$ for all $j = 1, \dots, J$ and $I = -40$. We choose this value of I to provide households with the income, net of child care expenses, that guarantees that their consumption is at least equal to the minimum daily living expenses observed in the data for households with a young child.

On the supply side, we must choose the number of potential entrants, K . In this choice we face a trade-off: the higher the K , the more heterogeneity of firms – but also the higher the computational cost of solving the two-stage game on the supply side. We strike a balance by setting $K = 15$. Since a potential entrant is defined by a pair of marginal costs (c^L, c^H) , we draw these pairs from a joint-lognormal distribution with a mean vector equal to (μ_{c^L}, μ_{c^L}) and a variance-covariance matrix equal to a diagonal matrix with parameters $(\sigma_{c^L}^2, \sigma_{c^L}^2)$. All potential entrants face a fixed cost F to enter the market. We set $F = 204.56$ based on industry cost reports (see Mitchell and Stoney (2012) and footnote 23 for more details).

Recall that a market configuration, \mathfrak{S} , is the set of firms operating in the market as well as their qualities. Given the number of potential entrants and possible qualities, there are clearly many possible market configurations. Note that, while there are K potential entrants, in a given market configuration only some of them might actually enter. In addition, two configurations might have the same number of actual entrants yet differ in the identity of those entrants.

For a given parameter point, we prepare for the computation of the equilibrium by finding the profit function, $\pi(\mathbf{p}; \mathfrak{S})$. For each possible market configuration \mathfrak{S} , this function gives the maximum profits for each firm. To compute these profits, for each market configuration we find the equilibrium of the Bertrand sub-game of price competition among the corresponding firms. In particular, we compute the vector of optimal prices $\mathbf{p}(\mathfrak{S})$ using a fixed-point algorithm on the first-order conditions defined in equation (14). Thus, the profit function gives the equilibrium of *each* sub-game of the second stage of the supply-side model. Note that this is *not* the equilibrium of the model, but rather the equilibrium of each sub-game for a given market configuration. In order to find the equilibrium market configuration, \mathfrak{S}^* ,

we must solve the whole game. Appendix B below explains the equilibrium computation.

B Computation of the Equilibrium

To compute the equilibrium for a given parameter point, we proceed as follows:

1. For each of the J household types, we draw a pair of wage offers (w_a, w_b) from a bivariate lognormal distribution with parameters $\mu_{w_a}, \mu_{w_b}, \sigma_{w_a}, \sigma_{w_b}$ and ρ .
2. We generate K random draws of (c_k^L, c_k^H) from a bivariate lognormal distribution with mean vector $\boldsymbol{\mu}_c$ and variance matrix $\boldsymbol{\Sigma}_c$. Each draw represents a potential entrant.
3. We establish the entry order of the K potential entrants based on their profitability (and hence efficiency), as described in Appendix C. While this step establishes the order of entry, it does not specify which firms enter (as some potential entrants might not enter), or which quality they offer.⁴³
4. Recall that, for every market configuration, the profit function $\pi(\mathbf{p}; \mathfrak{S})$ gives the optimal payoffs for each firm, and hence tells the quality (low or high) that maximizes profits for each firm under a given market configuration. Given the order of entry from step 3, and the optimal price and quality for every firm given by $\pi(\mathbf{p}; \mathfrak{S})$, in this step we solve the first (entry) stage of the supply side model by backward induction, as each firm decides whether to enter the market at its optimal quality and price, or whether not to enter. The solution of the first stage yields the equilibrium market configuration, \mathfrak{S}^* .

With K potential entrants, each facing three choices (not entering, entering with high quality, and entering with low quality), there are 3^K configurations of available providers and child care qualities. Since the equilibrium of the entry stage is found by backward induction, in principle all 3^K scenarios of the competition stage have to be computed. This is computationally burdensome for $K > 10$.

To lessen the computational burden, we apply a refinement to limit the number of configurations. In this refinement, a potential entrant with order of entry k does not consider entering if potential entrant with order of entry $k - 1$ decides not to enter. This is because the order of entry reflects firm efficiency. As a result, if a more efficient firm decides not to enter, the less efficient firm should not enter either. We can then eliminate all configurations in which potential entrant k enters the market while the $k - 1$ previous providers do not. This turns the entry stage into a subgame in which a potential entrant enters only if the previous ones have entered, and the number of all possible configurations falls to 2^K .

C Definition of firms' entry order

This appendix describes the algorithm used to define the entry order of potential entrants in the first (entry) stage of the model. The algorithm includes the following steps:

⁴³In Stackelberg games, players move sequentially. In our Stackelberg entry game, we define the sequence of entry based on firm profitability, which is a reflection of firm efficiency.

1. From the set of K potential entrants, find the pair $\{r, k\}$ that have the minimum realizations of c_k^L and c_r^H , respectively. These are the most efficient potential entrants for low and high quality, respectively.
2. From the expected profit function $\pi(\mathbf{p}; \mathfrak{S})$, retrieve the optimal profits that r and k would obtain if each one were to be the sole provider in the market (i.e., a monopolist offering the quality level for which it is more efficient than any other firm).
3. Compare the profits from step 2 for r and k . Of these two firms, the one with the higher profits is the first entrant. In other words, the first entrant is the firm that would be most profitable as a monopolist.
4. Repeat step 1 for the remaining $K - 1$ providers. This will yield a new pair $\{r, k\}$.
5. From the expected profit function $\pi(\mathbf{p}; \mathfrak{S})$, retrieve the optimal payoffs for the new r and k as if each of them, separately, were competing with the first entrant, given the first entrant's quality. The one with higher profits is the second entrant. In other words, the second entrant is the firm that can compete most effectively with the first entrant.
6. Repeat steps 4-5, until all potential entrants are assigned a position in the entry order.

D Data Appendix

The data we use in estimation are extracted from the US Early Childhood Longitudinal Study birth cohort (ECLS-B). The study follows, from birth through kindergarten, a nationally representative sample of children born in 2001. The data are collected over several waves to capture child development. We focus on the second wave, collected between January and December 2003, when the children were two years old.

As mentioned in Section 3, we arrive at the estimation sample by imposing three main restrictions on the full dataset: we focus on two-parent families with one 2-year old child; we limit work status to working full-time or not working; and we consider only parental or center-based care. In addition, we eliminate observations with hourly earnings in the top or bottom 0.5 percent of the distribution, conditional on gender and education (high school / college).

In Table D.1 we report how these restrictions affect our estimation sample as they are progressively applied. Per the first restriction (second column), we keep households with three members - mother, father, and a two-year old child. We do not drop households that go on to have more children after our sample period, but we drop households in which the two-year old child has other siblings (biologically or through adoption) during our sample period, or that include additional adults.

Per the second restriction (third column), from the previous set we select the households where the father is working full-time and the mother is either working full-time or not working. Full-time is defined as working for pay for no less than 40 hours a week. We eliminate all households where at least one parent does shift work, since shift workers have distinct child

care needs. We eliminate a few (eight) observations for which child care use seems measured with error, as both parents work full time (not from home) yet do not use outside child care.

Per the third restriction (fourth column), from the previous set we focus on parents that use either parental care or a combination of parental care and paid center-based care. Note that all parents in the data provide at least some parental care.

The estimation sample is composed of families that are, on average, more educated than those in the full dataset. As a result, these families earn, on average, higher wages, and mothers have a higher participation rate. The child's cognitive score is also higher than in the full dataset, by about 2.6 percent.

Table D.1: Estimation Sample Comparison

Sample: Restrictions:	<u>Raw data</u>	Only one 2 year old	Father working FT	<u>Estimation Sample</u> Only parental or center care
<u>Child's cognitive Score</u>				
Average	125.49	127.81	128.88	128.83
Standard Deviation	10.98	11.53	11.67	11.3
Missing	0.088	0.096	0.073	0.074
<u>Parents' Labor Market:</u>				
<u>Mothers</u>				
college completed or more	0.27	0.41	0.52	0.49
Proportion Working	0.61	0.65	0.38	0.28
Wage: Average	15.63	17.45	20.84	20.81
Wage: Standard Deviation	11.95	12.00	12.14	13.42
<u>Fathers</u>				
college completed or more	0.33	0.39	0.51	0.52
Proportion Working	0.95	0.96	1.00	1.00
Wage: Average	20.91	20.69	23.5	23.87
Wage: Standard Deviation	17.16	17.6	16.41	17.17
<u>Sample size</u>				
Number of Households	9,650	1,650	550	450

NOTE: Samples extracted from the second wave of ECLS-B for children born in 2001. Moving left to right, we progressively add the restrictions leading to the estimation sample. Sample sizes are rounded to the nearest 50 in order to comply with ECLS-B confidentiality rules.