Household Economies of Scale: Benefits for Age-grouped Children?

Kelly M. Jones*

August 29, 2008

Abstract

Estimates of economies of scale in household size and composition have long been used to adjust calculations of household welfare. It is suggested here that there are similar economies found in other aspects of household structure, such as the age distribution of the children within it. Age-grouped children may benefit from economies of scale that act as price-effects for time-intensive child rearing inputs. As the number of children benefiting at once from the time investment increases, the cost per child decreases. Such price effects may significantly increase demand for child-quality inputs, particularly in poor countries. An empirical test for such economies examines the effect of age-cohort size on demand for childhood immunizations in Senegal. Results offer evidence that economies of scale do arise from clumpy age distribution.

*University of California Berkeley. The author thanks Jeremy Magruder, Alain de Janvry, Maximilian Auffhammer, Jennifer Ifft, Abdoulaye Sy, Clair Null, and the participants of the ARE Development Workshop for helpful comments and suggestions.

1 Introduction

In 1948, Frank and Earnestine Gilbreth's book *Cheaper by the Dozen* popularized the economic concept of economies of scale in household size. When anyone asked why the Gilbreth's had *so many* children (twelve!), Frank always replied, "Well, they come cheaper by the dozen, you know!" It is easy enough to observe in daily life, that public goods such as housing or an automobile are cheaper to each person when shared. Cooking for two is certain to save on food and time costs compared to each cooking for oneself.

Consumption benefits from public goods, bulk purchasing and time division within a household have long interested economists. Over a century of literature has been dedicated to assessing the role that household size and composition play in consumption decisions. Attempts to quantify economies of scale and estimate equivalency scales across households are numerous. Such measures are useful in both research and policy for the calculation of household welfare and poverty lines. However, all measures to date have accounted only for household size and perhaps composition of men, women and children.

In this paper I propose that other characteristics of household structure may also offer economies similar to those arising from household size and composition. In particular, this work focuses on the age distribution of children within a household. The question of interest is how the size of one's age-cohort affects demands for child-quality inputs, holding total number of children constant. Because certain child inputs can be time-intensive on the part of adults, providing these for multiple children at once may reduce the marginal cost. A simple model predicts that if economies in inputs are gained by age-grouping, children with larger age-cohorts, ceteris paribus, should exhibit greater demand for such inputs.

Because demand for child-quality inputs can be highly price-elastic in poor countries, a reduction in marginal cost could be quite effective in increasing human capital investments. In this paper I present an empirical test for evidence of economies of scale in age-cohorts using data on households and children in Senegal. As an example of childhood inputs, I examine the relationship between immunization status and a child's age-cohort size.

To estimate the effect of cohort size holding all else constant, I compare children to others within the same household via household-fixed effects. However, many households in Senegal contain several mothers of young children and so unobserved mother characteristics may still enter the error term, despite the household-fixed effect. While this seems to complicate the analysis, at the same time it offers a solution. I exploit these co-residencies and use the number of *non-sibling* children in the household close to one's age as a measure of the child's age-cohort size. These non-sibling age-mates should offer the same economies of scale within the household as siblings. However, use of this measure rather than siblings breaks the endogenous connection between birthing and input decisions. Finally, I control for observable differences between children that might affect parental input decisions, such as sex and birth order.

The identified effects of cohort size are positive and statistically different from zero. Within a household, children with more coresidents their own age are more likely to receive immunizations than children with a smaller age-cohort. The coefficients suggest that each age-mate increases the likelihood of vaccination by 1 to 2 percentage points. Such an increase is non-negligible given the current national immunization rate of 73%. These estimates are robust to a variety of sub-sample selections and definitions of age-mate. The increased demand resulting from additional age-mates suggests that economies of scale in age-cohorts are creating a price-like effect for child immunizations. As modeled, such economies may arise from any child-quality input for which non-excludable adult-time is required for provision.

The remainder of the paper is organized as follows. Sections (2) and (3) review past economic studies of household economies of scale and comment on the contribution of this paper. A basic model of the mechanism of scale economies is presented in section (4). The empirical analysis follows in section (5), including a discussion of the data, methodological challenges and estimations. Sub-sections (5.3) and (5.4) present the results and consider alternative explanations for the findings, and section (6) concludes that there is evidence for economies of scale in age distributions of children.

2 Previous Study of Household Composition Effects

The primary focus of early work in this field was the effect of household demographics on expenditure choices. Estimates of such effects date back to the late 19th century. Earnst Engel [1895] was one of the earliest researchers to calculate equivalency weights for members of a household, relative to a numeraire member. Such weights, calculated by Engel and others, were generally based on physiological, rather than economic, reasoning.

Decades later, authors such as Prais [1953] and Barten [1964] were interested in whether, for a given household income, an increase in the size of the household reduces the per-capita consumption.[1] Barten's model, used frequently in the literature (see Muellbauer [1977] and Pollak and Wales [1981]), predicts that a change in family composition, holding total income constant, will change the household's consumption choices across goods. He likened these "family effects" to price effects. Barten, Prais, and others suggested that larger families are *worse off*, given the same total income as a smaller family.

Beginning in the 1980's, thinking moved toward the concept that for a given *per-capita* income, larger households were *better off*. This intellectual advance was based on the concept of household economies of scale. Empirical evidence of this followed; Lanjouw and Ravallion [1995] found that accounting for such economies in household size provides evidence against the commonly held view that, in developing countries, larger families are generally poorer.

Deaton and Paxson [1998] proposed formalizing the mechanism of scale economies with a model based on the division of goods into public or private. They showed that as household size increases, holding per-capita income constant, public goods become effectively cheaper. This creates both substitution and income effects; yet for inelastic private consumption goods, they argue that the income effect will dominate. Thus in the presence of economies of scale, an increase in household size, ceteris paribus, will increase per-capita consumption of such goods. However, as pointed out by Gan and Vernon [2003], in reality many goods are neither entirely public nor private. Most goods have some aspects of both and are better described as *relatively* more public than one and more private than another.

[1] Much of this work was inspired by the implementation of family tax credits in the U.K. during this time and the resetting of the U.S. poverty line to reflect household size.

A more realistic model of economies of scale was proposed much earlier in an article by Lazear and Michael [1980] published in *American Economic Review.* In their theory, which underlies much of the recent research on the subject, the chosen measure of welfare is not market goods purchased, but rather the "service flows" resulting from such goods. In this way, a household's welfare is determined not only by income but also by its "environment," which determines the service flow resulting from any given market good. For example, it is not simply the ownership of an automobile that is valuable, but the service flow of using it. Clearly such service flow is also impacted by the number (and needs) of other individuals sharing use of the auto. In their article, Lazear and Michael interpret the "environment" parameter simply as the household composition.

3 Present Study of Household Structure Effects

As summarized in the previous section, economic thinking on the topic of household composition has progressed from the idea that for at given total income, larger households are at a disadvantage, to the idea that for a given per-capita income, large households actually experience a welfare advantage. Along the way, some thought has also been given to the proper adult-equivalency for counting children within a household. However, such considerations of household demographics and economies of scale have not ventured beyond simple measures of household composition; that is, counts of total household members as well as other composition factors such as members within given age-sex categories.

I argue here that other factors of structure surely impact many household decisions, consumption and otherwise. This study focuses on the the influence of one particular aspect of household structure, the age distribution of the children. In this work, I define an *age-cohort*$_m$ of a child as the other children in the household within m months of a child's age. Further, I will refer to other children within one's age-cohort as *age-mates* and to those with age-mates as *age-grouped* and those without as *age-isolated*. The primary question of interest is whether there are age-cohort economies of scale for age-specific child inputs. That is, holding all else about a household constant, are children in larger age-cohorts at an advantage in receiving costly inputs?

One may conceive of several scenarios in which the size of one's age-cohort may have price-like effects on inputs for a child. The primary reason for the effects is generally the parental time-intensive nature of providing such inputs. For example, one may think of reading to a child as a quality input that is costly of adult time. Clearly reading to two or three children at once has a much lower per-child cost. Enrichment activities such as Saturday soccer games present a similar example.

In the context of poor countries, the case becomes even more clear. Because inputs such as nutrition, basic schooling and health care are often not affordable, even small price effects could significantly increase demand. For example, in many African households, girl children take care of the home, allowing the mother to work in the farm. A girl with sisters close in age may be offset at home and thus may be more likely to attend school than a girl with sisters much older or younger than herself. Other inputs such as taking children to a health clinic for immunizations or treatment are also quite costly of adult time. Long distances and poor transportation can make travel to service sites quite burdensome. In such cases, taking several children at once can create considerable savings on the per-child cost of the input. Thus, for a child with age-mates the cost may be half or less of that for an age-isolated child.

The primary contribution of this work is a look beyond household size and basic composition to another factor of household structure as a source of household economies of scale. Understanding the benefits or disadvantages of certain aspects of household structure is important for welfare analysis and social policy. A multitude of work has attempted to calculate correct equivalency scales based on household composition. However, if scale economies are also derived from other aspects of structure, surely these aspects should be included in such calculations as well.

Additionally, the vast majority of studies have used measures of food expenditures as a share of nominal income as evidence of economies of scale. Certainly these are not the only indicators of a household's welfare. Because I am focusing on child rearing in this analysis, the measures of welfare used here are those that indicate investments in child quality. Particularly in poor countries, it seems that price effects that render more affordable investments in future human capital are of utmost importance.

4 Modeling Household Economies of Scale

In this section, I first describe a general model of scale economies within a household and then apply it to the particular type of economies studied here. I then lay out the predictions of the model to be tested using household data.

Following Lazear and Michael [1980], I define a service flow, S_1, as the benefit derived from a costly good, X_1. The ratio of benefit to good is determined by the environment of the individual, so we define

$$\alpha_1 \equiv \frac{S_1}{X_1}$$

where α_1 is an environment parameter. Then the cost of receiving the benefit S_1 is

$$P_{S_1} = \frac{P_1}{\alpha_1} \tag{1}$$

where P_1 is the market cost of good X_1 .

Therefore, an individual's demand for S_1 is

$$S_1 = d\left(\frac{P_1}{\alpha_1}, Y\right) \equiv g(P_1, Y)$$

where Y is nominal income, and the derived demand for X_1 is

$$X_1 = \frac{g(P_1, Y)}{\alpha_1} \equiv h(P_1, Y).$$

Although Lazear and Michael interpreted α_1 as an indicator of household composition only, the same model can be applied to household structure. This allows one to formalize the concept of economies of scale within age-cohorts for important childhood inputs such as schooling, visits to a health clinic, or immunizations. As discussed in the previous section, a significant cost of inputs of this type is the contribution of adult time.

I let S_{iv} represent child i's receipt of input v, and X_{iv} represent the adult time required to provide the input. As before, we define a child's environment parameter

$$\alpha_{iv} \equiv \frac{S_{iv}}{X_{iv}}$$

as the determining factor of how much S_{iv} is received per X_{iv}. As illustrated in the examples above, α_{iv} can best be thought of as the size of the child's age-cohort within her household. This environment parameter is indexed also by v, as the particular input of interest would define the relevant age-distance for a cohort. For example, considering attendance at primary school, age-mates within 6 years of oneself may be relevant. However, for early childhood inputs such as immunizations, only those within 2 years oneself would create the price effects considered here.

Revisiting equation (1), we add a term to capture the entirely private costs of inputs that do not change with age-cohort size (e.g. the actual cost of the school fees and a uniform)

$$P_{S_{iv}} = \frac{P_v}{\alpha_{iv}} + \varepsilon_v. \tag{2}$$

Despite these additional costs, $P_{S_{iv}}$ is still decreasing in α_{iv} (under constant labor market prices) and thus S_{iv} is increasing in α_{iv}.

Thus if age-cohort scale economies exist, we expect to observe the following effect:

$$\frac{\partial S_{iv}}{\partial \alpha_{iv}} > 0 \tag{3}$$

5 Empirical Analysis

Given the likely importance of these price effects in poor countries, the empirical test of relation (3) presented here is set in the country of Senegal. While the theory is applicable to numerous critical childhood inputs, for simplicity, the test focuses on a single, representative input: that of early childhood vaccinations. The Senegalese National Immunization Schedule recommends the timing of ten primary immunizations as shown in Table 1 [WHO, 2007].

5.1 Data

I employ data from the Demographic and Health Survey (DHS) for Senegal collected in early 2005, which is a nationally representative sample of 7,412 households. The DHS interviews all women aged 15-49 in the sampled households on topics relevant to fertility, reproductive health, marital relations, and childhood health and nutrition. This sample includes 14,602 women, 65% of whom are mothers.

The survey records whether each child born since 2000 did or did not receive each vaccine. Therefore, the full sample for this analysis is 11,607 children, born to 7,306 mothers in 4,845 households.

A summary of the coverage of immunizations is also presented in Table 1. Overall coverage of immunizations is mediocre, 73% on average, and earlier vaccines are more commonly received. Vaccines for which a child is not yet due are excluded from the mean, giving rise to the slight decline in sample size for vaccines due at older ages.

Table 1: Senegalese National Immunization Schedule & Coverage

Vaccine Dose	Due at age	Of children due, % that have received	Sample Size
Tuberculosis	Birth	87%	10,045
Oral Polio Vaccine 0	Birth	46%	10,035
Diphtheria, Pertussis, Tetanus 1	6 weeks	86%	9,951
Oral Polio Vaccine 1	6 weeks	87%	9,967
Diphtheria, Pertussis, Tetanus 2	10 weeks	80%	9,859
Oral Polio Vaccine 2	10 weeks	78%	9,871
Diphtheria, Pertussis, Tetanus 3	14 weeks	68%	9,777
Oral Polio Vaccine 3	14 weeks	62%	9,789
Measles	9 months	67%	9,335
Yellow Fever	9 months	66%	9,290
All Vaccines		73%	97,919

As a preliminary check of the validity of the model, a subsample is examined for evidence that age-grouped children are often taken as a group to receive immunizations. For a portion of the full sample, the enumerator was able to view the healthcard that lists the immunization history of a child. For this healthcard sub-sample, which includes 56% of children, the data contain the date on which each immunization was received. With this additional information we can see whether the vaccine was received when it was due (or later). As shown in Table 2, over 10% of households with age-grouped children have multiple children receiving immunizations on the same day. This differs significantly from households that may have multiple young children but do not have any that are age-grouped within 2 years of each other.

Further, it is notable that immunizations are scheduled at the earliest age when it is safe for a child to receive it. Therefore, a vaccine is either received when it is due, or it is received late, as it cannot be administered early. Although all immunizations are due by age 9 months, the data reveal that many children receive them late – even up to age 48 months. We would expect that taking children as a group would therefore decrease the likelihood of receiving immunizations exactly when they are due. Thus conditional on receiving a vaccine, the likelihood of receiving it late should be higher for age-grouped children. Evidence of this is shown in the second row of Table 2: of vaccines received in households with age-grouped children, 47% were received late. Again, this is statistically different (at the 1% level) from households that do not have age-grouped children, suggesting that age-grouped children are more likely to be taken as a group for immunizations.

5.2 Econometric Test

In order to more formally answer the question of interest, I test for evidence of relation (3),

$$\frac{\partial S_{iv}}{\partial \alpha_i} > 0.$$

Table 2: Evidence of age-grouping for immunizations

	Households with(out) age-grouped children		
	Without	With	t-statistic for difference
Percent with kids receiving vaccines on the same day	0.13%	11.64%	17.35
Percent of received vaccines that were late	44.60%	47.0%	4.64

That is, what is the impact a child's cohort size on the amount of time-intensive inputs she receives? As noted above, the input used for this test is a child's immunization.

5.2.1 Methodological Challenges

Several challenges exist in ascertaining the sign of $\frac{\partial S_{iv}}{\partial \alpha_i}$. In order to be sure that it is a difference in α_i causing a difference in S_{iv}, one must rule out any other potential causes for such differences. To begin, we note that one cannot simply examine the correlation between cohort size and the number of immunizations received. Age grouped children may, on average, come from households that are different in important ways from those that house age isolated children. In particular they are likely to be in households that have more children in total. Further, household characteristics such as wealth, age and education of the household head, religion, and a host of unobservable characteristics may influence decisions regarding health inputs as well as those regarding reproduction. In order to avoid any bias resulting from the correlation of these excluded characteristics with the variables of interest, I employ a within-estimator using household fixed effects to compare across children within the same household. In combination with an indicator for a child's birth order, this holds constant the total number of children in the household when the child is due for immunizations, as well as time-invariant household characteristics.

However, since nearly 50% of the households in the sample include multiple mothers of young children, one may still be concerned that unobservable differences between mothers within a household could influence both birth timing and immunization decisions.[2] For example, if a mother is by nature very concerned for her children's health, she is likely to employ adequate birth spacing, resulting in age-isolated children that are also likely to be vaccinated. This would cause a downward bias in the estimation of the relation of interest. Mother-fixed effects could remedy this problem but are unfortunately not feasible in this sample. Variation in age-cohort size across a mother's children would require her to have at least 3 children in the sample. Since vaccination information is only available for children born within 5 years of the survey, less than 10% of mothers have 3 or more children in the sample.

Mother-fixed effects being unavailable, one might consider the use of an instrumental variable to

[2] In this analysis (as in the data), a family is defined as a mother and her children. A household is defined as members who coreside – generally extended families, but not necessarily.

Table 3: Characteristics of the Sample

	Proportion or Mean	Std. Dev	Min	Max
Households				
Number of children under age five	2.34	1.52	1	15
Have age-grouped children	50%			
Have multiple mothers	47%			
Has variation across children in NSC age-cohort size	33%			
N=4,845				
Mothers				
Age	29.4	7.44	15	49
Number of children	3.9	2.62	1	14
Children under five	1.6	0.66	1	5
Have coresident mothers	64%			
Variation across children in NSC age-cohort size	24%			
In polygamous union	34%			
Getting to a clinic is "big problem"	40%			
N=7,306				
Children Under Five				
Percent of vaccines received	0.69	0.35	0	1
Healthcard data recorded	52%			
In HH with multiple mothers	64%			
NSC age-cohort size	0.98	1.41	0	14
N=11,607				

avoid endogeneity in the estimation. Yet an instrument that predicts a child's age-grouping while being independent of her likelihood to receive vaccinations is elusive given the close relationship between such decisions. As an alternative, I exploit the prevalence of coresident families in Senegal, allowing me to exclude a child's siblings from the calculation of her age-cohort size. The remainder in her cohort will be referred to as non-sibling coresident children (NSC). Recall that it is the increase in one's age-cohort size that is expected to generate the price (and thus demand) effects. Households operating as a unit would likely take all young children to the clinic together, so that non-sibling coresident age-mates would offer a benefit similar to that of age-mate siblings.[3] Focusing on NSC's creates a measure of that environmental parameter that is free of the endogenous relationship between reproductive and input decisions for different children.

After controlling for any potentially endogenous household factors with fixed effects and evading unobserved mother characteristics by focusing on NSCs, one may still be concerned about important differences across children within a household. In order to capture any differentials in likelihood of immunization based on child characteristics, I also include the child's gender, sibling birth order, and household birth order in the equation.

[3] For evidence of cooperation in coresident families, see Foster [2004], which states that "the health status of coresident family units are viewed as complementary to the health of one's own family".

5.2.2 Estimations

In estimating the effect of age-cohort size on likelihood of receiving immunizations, I first let $m = 24$. That is, I define the relevant age-cohort for vaccinations as the number of other children within 24 months of one's age. This is based on the immunization schedule as shown in Table 1, combined with information from the data suggesting that children continue to receive vaccines up to age 48 months. However, 99% of vaccines received are completed by age 24 months.[4]

As discussed in the previous section, the measure of age-cohort is NSC_{ij}, the number of non-sibling coresident children within 24 months of one's age. The primary equation of interest is

$$S_{ijv} = \beta NSC_{ij} + Z_{ij} + \gamma_j + \gamma_v + \varepsilon_{ijv} \tag{4}$$

where S_{ijv} indicates whether immunization v was received by child i in household j, and is predicted by NSC_{ij}, controlling for Z_{ij}, a vector of child-specific characteristics as discussed above, a household-fixed term γ_j, a vaccine-fixed term γ_v, plus a random, mean-zero error term. The central identification assumption of this approach is the following: controlling for anything common across mothers in a household by the use of household-fixed effects, the immunization decisions of a mother are independent of the spacing decisions of her coresident mothers.[5]

The theory predicts that β will be positive and significantly different from zero, indicating that a child's likelihood for receiving a vaccination is increasing in the size of her age-cohort. Such a result would offer evidence that households do experience economies in time-intensive child inputs when children are age-grouped into larger cohorts.

5.3 Findings

5.3.1 Primary Results

The characteristics of the sample regarding coresident mothers and children are detailed in Table 3. Forty-seven percent of households have multiple mothers of children under age 5, and 33% of households exhibit variation across children in age-cohort size. While 64% of women have coresident mothers, only 34% are in polygamous unions, suggesting that many coresident mothers are likely sisters, sisters-in-law or cousins, etc. The average number of NSC age-mates for children in this sample is 0.98.

Columns (1) and (2) of Table 4 show OLS estimations of equation (4) for $m = 24$; the first excludes child characteristics. Both columns show a positive and statistically significant impact of age-cohort size on immunization status based on robust standard errors clustered at the household level. The difference in the estimates of β between the two columns is slight. The decrease in the point estimate resulting from inclusion of these characteristics likely results from the fact that children born later in a household (thus with higher values for household birth order) are statistically less likely to receive vaccinations. Because these children are also more likely to be age-isolated[6], column (1) is slightly

[4]Checks for robustness include varying the size of m between 12 and 48 months.

[5]Note that for single-family households, the household-fixed effect is, in essence, a mother-fixed effect. As a specification check, these, as well has polygamous households, are excluded from the sample and it does not change the results.

[6]Estimating within households, I find that for each step toward the end of the birth order, the average number of NSC agemates is decreased by 0.8.

Table 4: Effect of Age-cohort Size on Immunization Status

	Dependent Variable: (Y/N) Vaccine was received For all columns, $m = 24$			
	(1)	(2)	(3)	(4)
NSC age-cohort size	0.017***	0.012***	0.018***	
	(0.004)	(0.004)	(0.006)	
NSC age-cohort size^2			-0.001	
			(0.001)	
Number of NSCs not in age-cohort				-0.009**
				(0.004)
Male		0.019***	0.019***	0.019***
		(0.006)	(0.006)	(0.006)
Birth Order		-0.001	-0.001	-0.007***
		(0.002)	(0.002)	(0.003)
HH birth order		-0.025***	-0.025***	-0.024***
		(0.002)	(0.002)	(0.002)
Constant	0.640***	0.839***	0.840***	0.889***
	(0.006)	(0.021)	(0.021)	(0.024)
Vaccine fixed effect	Yes	Yes	Yes	Yes
Household fixed effect	Yes	Yes	Yes	Yes
Observations	92791	92791	92791	92791
R-squared	0.38	0.39	0.39	0.39

Notes: Robust standard errors are in parentheses, clustered at the household level. Astrices indicate the standard levels of significance: 10%, 5% and 1%. "NSC's not in age-cohort" indicates a coresident under age 13 yrs but not within 24 months of one's age. All columns are estimated at the child-vaccine level and vaccines for which a child is not yet due are excluded.

upward biased by the correlation of this excluded factor.

The estimation in column (2), controlling for relevant child characteristics, suggests that each additional NSC age-mate increases the chance a child will be vaccinated by 1.2 percentage points (for $m = 24$). Based on current immunization rates above 60%, this is not a large increase, however it is statistically different from zero at the 1% level. Additionally, it is important to note that increasing rates by just 1.2 percentage points would result in an additional 47,000 vaccines received per year in Senegal.[7] Further, this point estimate is similar in magnitude and sign to the differential treatment boys receive in terms of health inputs. That is, for a girl, having one additional agemate nearly makes up the gender gap in chance of vaccination.[8]

To check for potential non-linearities in this relationship, column (3) includes a quadratic term. That the coefficient on the quadratic term is indistinguishable from zero suggests that the effect is linear in nature. Further, the point estimate on the primary variable of interest increases only slightly, and is not statistically different (at the 1% level) from the initial estimate of 1.2.

For comparability, column (4) explores the impact of non-sibling coresident children (under age 13) that are *not* within 24 months of one's age. The results show that the presence in the household of children outside one's own age-cohort in fact has a negative impact on immunization status.

5.3.2 Specification checks

A potential concern regarding the estimations shown thus far is that the small standard errors may be a result of the large sample sizes due to estimation at the child-vaccine level. However, because the standard errors are clustered at the household level in all estimation, this should not be a concern. Nevertheless, as an alternative specification, column (1) of Table 5 displays the estimation at the child level. Here, the dependent variable is the percent of vaccines (for which a child is due) that he or she has received. The sign, magnitude and significance are stable for the coefficient on number of non-sibling coresident age-mates.

To test the sensitivity of the results to the choice of age-cohort cut-off, m, columns (2) - (4) of Table 5 estimate equation (4) with m set to 12, 36, and 48 months. In all cases, the coefficient on cohort size remains positive and significantly different from zero. Of note, as the window is widened, that is as the age-cohort includes children farther from one's own age, the magnitude of the effect diminishes. Figure 1 plots the estimated coefficients against the range of values for m. We see that the closer to one's age an age-mate is, the greater impact he will have on a child's chance of immunization.

In Table 6, I test the sensitivity of the results to specific subgroups within the sample. To allay any concerns about lack of sufficient variation in the independent variable of interest, column (1) re-estimates equation (4) including only observations from households that exhibit variation across children in cohort size. The coefficient estimate and its standard error are remarkably stable, indicating that each additional agemate increases a child's likelihood of immunization by 1.3 percentage points.

[7] Based on a population of approximately 12 million in 2006, and a birth rate of 32.78/1,000 people, 3.934 million vaccines are due per year in Senegal.

[8] In comparison to columns (1) and (2), equations are also estimated using the sibling age-cohort size as the independent variable of interest. As discussed above, unobserved maternal characteristics that are correlated with short birth spacing and low concern for immunizations are likely to make the use of siblings endogenous. One would expect such endogeneity to cause a downward bias on the estimates. As expected, the estimated parameters are significantly lower than those in table 4 and are, in fact, indistinguishable from zero. These results are not presented here but are available

Table 5: Specification Checks

	Child-level % of Vaccines	Child-vaccine level (Y/N) Vaccine Received		
	m=24	m=12	m=36	m=48
	(1)	(2)	(3)	(4)
NSC age-cohort size	0.013**	0.017***	0.011***	0.006*
	(0.006)	(0.005)	(0.003)	(0.003)
Male	0.004	0.018***	0.019***	0.019***
	(0.008)	(0.006)	(0.006)	(0.006)
Birth Order	0.002	-0.002	-0.001	-0.001
	(0.002)	(0.002)	(0.002)	(0.002)
HH birth order	0.025***	-0.026***	-0.023***	-0.023***
	(0.003)	(0.002)	(0.002)	(0.003)
Constant	0.429***	0.856***	0.821***	0.829***
	(0.027)	(0.020)	(0.024)	(0.025)
Vaccine fixed effect	No	Yes	Yes	Yes
Household fixed effect	Yes	Yes	Yes	Yes
Observations	11502	92791	92791	92791
R-squared	0.52	0.39	0.39	0.38

Notes: Robust standard errors are in parentheses, clustered at the household level. Astrices indicate the standard levels of significance: 10%, 5% and 1%. Column 1 is estimated at the child-level and the dependent variable is share of vaccines received; columns 2-4 are estimated at the child-vaccine level and the dependent variable is (Y/N) the vaccine was received. In both cases, vaccines for which a child is not yet due are excluded.

Figure 1:

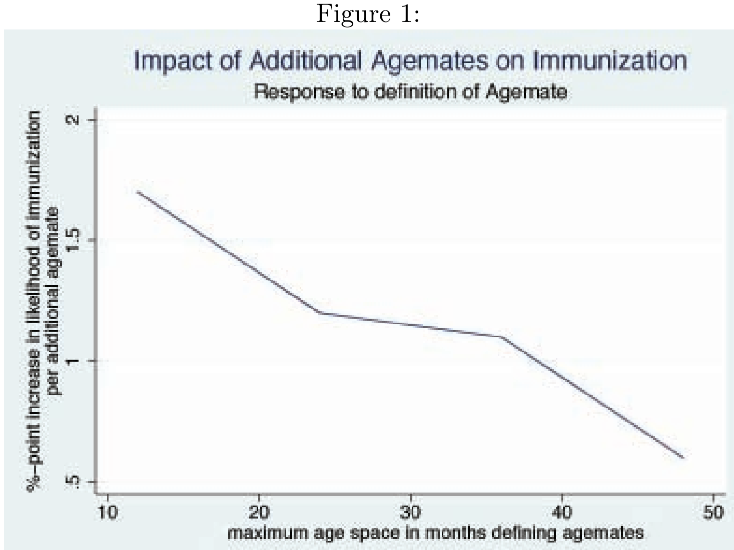

Table 6: Specification Checks & Tests for Alternative Explanations

	Dependent Variable: (Y/N) Vaccine was received			
	Specific Sub-Samples		Full Sample	
	HHs with variation	Excluding polygamous	Shocks last $\leq$ 12 mo	Shocks last $\leq$ 24 mo
	(1)	(2)	(3)	(4)
NSC age-cohort size	0.013***	0.012**		
	(0.004)	(0.005)		
"Uncorrelated" NSC age-cohort size			0.015***	0.017***
			(0.005)	(0.005)
Male	0.021***	0.028***	0.019***	0.019***
	(0.007)	(0.009)	(0.006)	(0.006)
Birth Order	-0.001	-0.001	-0.002	-0.002
	(0.002)	(0.003)	(0.002)	(0.002)
HH birth order	-0.022***	-0.025***	-0.028***	-0.028***
	(0.002)	(0.003)	(0.003)	(0.003)
Constant	0.859***	0.870***	0.869***	0.869***
	(0.028)	(0.033)	(0.021)	(0.020)
Vaccine fixed effects	Yes	Yes	Yes	Yes
Household fixed effect	Yes	Yes	Yes	Yes
Observations	50947	34356	92791	92791
R-squared	0.34	0.37	0.39	0.39

Notes: Robust standard errors are in parentheses, clustered at the household level. Astrices indicate the standard levels of significance: 10%, 5% and 1%. "Uncorrelated" indicates that conception of age-mates was sufficiently distant from potential shocks (see text). All columns are estimated at the child-vaccine level and vaccines for which a child is not yet due are excluded.

Finally, one must exclude any households in which the key identification assumption could be violated. It is noted that while many households include multiple mothers of young children, a non-negligible percentage of these include polygamous unions. In the case of polygamy, it may be the case that reproduction decisions of one mother are no longer independent of a coresident mother, if they share a husband. In a one-husband household, the household fixed effect would control for such interactions. However, with multiple men, one or more of whom have multiple wives, then the central assumption of equation (4) may be violated. Therefore, column (2) of Table 6 presents the same estimation excluding observations from households containing any polygamous union.[9] Again, the results of this estimation are virtually identical to those from the full sample.

5.4 Alternative explanations for results

Peer Effects

One might suggest that the results shown here merely reflect peer effects among coresident mothers rather than true economies of scale. For example, if a mother observes her coresident mother investing in immunizations, perhaps she is more likely to also do so.

I argue that for this to be the primary mechanism behind these results, two rather unbelievable assumptions must hold. First, the peer effects must only operate in the positive direction. That is, observing investment must induce investment, but observing non-investment must not induce non-investment. If peer effects work in both directions, which seems more likely, then they would offset each other and the presence of age-mates would have an ambiguous effect. Secondly, one must assume that a mother is influenced only by coresidents with children of the same age as her own. Because we found no effect from non-sibling coresidents outside one's age-cohort, one must assume that mothers with children of other ages do not exert peer effects. This assumption seems even less tenable, as young mothers may look to senior women in the household for child-rearing advice.

Immunization Issues

Other potential explanations for the results relate to the choice of immunizations as the representative input used for the empirical test. Primarily, one might consider the potential externalities of vaccinations to be confounding factors in the analysis. That is, perhaps children in larger age-cohorts are more likely to be immunized as a result of their greater exposure to other children.

As before, there are stringent assumptions one must make for this reason to explain the results. First, children must be more contagious to age-mates than to children outside their own age-cohorts. This is perhaps more possible for some diseases than others. However, a second strong assumption is that those in one's age-cohorts account for a non-negligible fraction of a child's exposure to disease. Even the largest cohorts in this data (10+) are small compared to the number of children to which a child is generally exposed. Furthermore, diseases such as yellow fever and diphtheria are transmitted

from the author upon request.

[9] Note that whether the women in a household reporting polygamous unions all share the same husband or not is not known to the researcher. The number of married males on the household roster is not sufficient information, as many men are absent from the home for work or with live with other families. Therefore, all polygamous households are excluded as a conservative check.

via environmental factors. It is unlikely that the externalities *within a household* of immunizations could produce the results presented here.

Additionally, there could be concern with the examination of immunizations in Senegal since there are often "vaccination days" sponsored by a national campaign designed to encourage demand for the service by providing it at a lower cost. In fact, nearly two-thirds of children in the sample received at least one immunization on such an occasion. However, it is important to note that these campaigns do not entail door-to-door service. Generally, a service site is set up in or near a major market at a distance of up to 10km or more from some households. The campaign likely reduces the ε-cost of immunizations from equation (2). However, it does not eliminate the adult time required to receive the service, which is the source of the economies as modeled here.

Time-varying household characteristics

The presence of multiple young children in households allows for the exploitation of household-fixed effects, which is a useful way of dealing with the host of unobservable characteristics of a household. However, there may be important household factors that change through time that are not being captured by the fixed effect. To the extent that a changing factor is related to both reproduction and immunization decisions, it could confound the analysis. For example, if a household experiences a negative income shock and as a result reduces or postpones reproduction, the shock is likely to also reduce investments in child quality such as the inputs discussed here. If such an effect is operating, it would create a correlation between large cohorts and greater inputs and would create an upward bias in the point estimates.

In considering whether this effect could be driving the results, I first note that the time span included in this analysis is five years. Therefore, any longer-term shocks (covering all five years) are not considered variation through time. Further, given the lag between conception decisions and a resulting birth and considering the specific window of time during which a child receives vaccinations, shocks that affect both would be rare.

To test more explicitly whether such shocks should be a concern, I begin with the assumption that these negative shocks last up to a maximum of 12 months, an assumption which is later relaxed. The diagrams shown in figure (2) illustrate this discussion. If child i is born at time 0, then her window of vaccinations is $[0, 24]$, where time is measured in months. The time window for the birth of her age-mates is $[-24, 24]$, resulting in a window for age-mate conception of $[-35, 15]$. As shown in figure 2.B, I impose a window for household shocks affecting immunizations that extends 12 months beyond the immunization window in either direction. If an age-mate is conceived outside of this shock window, then the decision of her birth must be uncorrelated to any shock affecting input decisions for child i. Figure 2.C shows the resulting window for conception of "uncorrelated" age-mates and the associated window for births of such.

By this thought experiment, one finds that any age-mates born 26 to 3 months before one's own birth should be exogenous regarding any household shocks of concern. Relaxing the assumption imposed on maximum length of shocks from 12 to 24 months, one proceeds similarly to select uncorrelated age-mates. Under this assumption, uncorrelated age-mates are born 26 to 15 months before the child of interest.

Figure 2: Selection of Age-Mates Uncorrelated to Household Shocks

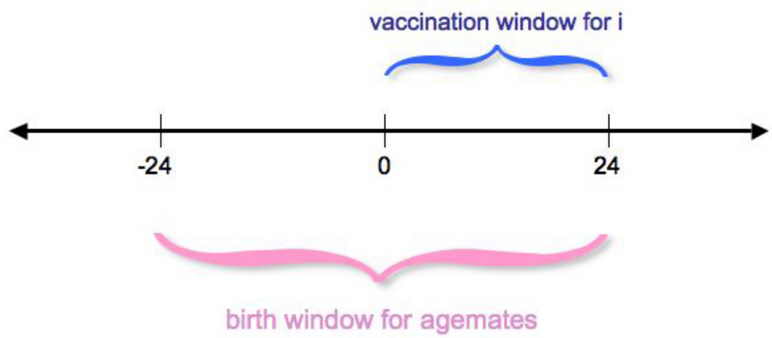

2.A

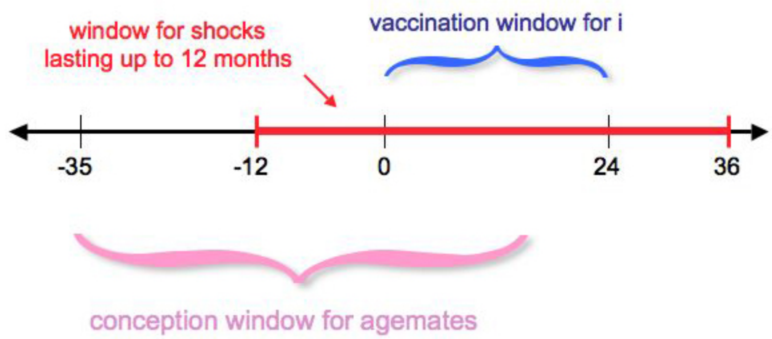

2.B

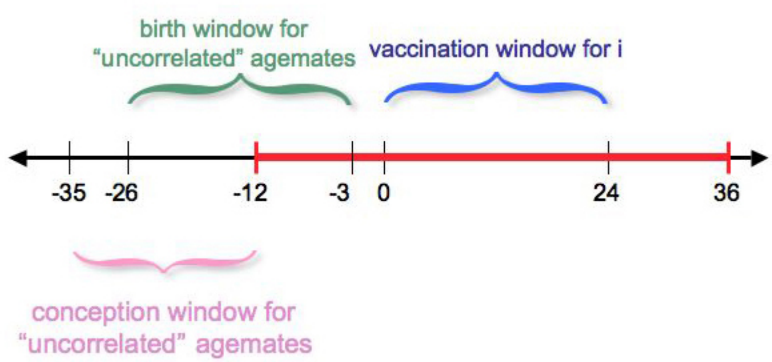

2.C

I re-estimate the previous specification, replacing NSC cohort size with the size of one's "uncorrelated" NSC cohort. If time-varying household characteristics are a factor here, one would expect that the parameter estimates would decrease as a result of excluding the previously upward biasing age-mates. The results however, shown in columns (3) and (4) of Table 6, show the opposite. Rather than decreasing, the point estimate actually increases under both the 12- and 24-month assumptions regarding shock length. This offers evidence against the importance of time-varying household characteristics in this analysis.[10]

6 Conclusions

The purpose of this study has been to investigate the potential for economies of scale arising from the age distribution of children within a household. While much research has examined economies of scale resulting from household size or composition, this is the first, to my knowledge, to seek evidence of scale economies arising from other aspects of household structure.

The economies examined here act as price effects, lowering the marginal cost of child-quality inputs for children in larger age-cohorts. Such economies are suggested to arise from the significant demands placed on adult time to provide these inputs. The proposed model predicts that these scale economies would result in a positive relationship between observed age-cohort size and inputs received.

In this analysis I have attempted to examine the causal relationship between these two observables by netting out as many other potentially influencing factors as possible. The use of non-sibling coresident age-mates as a measure of one's age-cohort eliminates problems of endogeneity arising from unobservable mother characteristics. Household fixed effects are used to compare children within a household, thus eliminating issues of household unobservables.

Using childhood immunizations as a representative input for the empirical test, I find that having a larger age-cohort results in a greater likelihood of receiving a vaccination, relative to a child in the same household who has a smaller age-cohort. These results are robust to the inclusion of child-specific characteristics, as well as the definition of age-cohort and various restrictions on the data sample.

The estimation suggests that each additional age-mate increases the chance of immunization by 1 to 2 percentage points. Unfortunately, in order to calculate the size of the price effect that is inducing this observed change in demand, one would need knowledge of the demand function for a vaccination. Lacking this, estimation from these results of a quantitative measure of the economies provided by the marginal age-mate is not possible. However, what can be said is that the induced demand increase of 1 to 2 percentage points is non-negligible. Applying such an increase to the current national immunization coverage rates would result in nearly 50,000 additional vaccines received per year. Considering the multiplicative public health benefits resulting from increases in coverage, the social benefit of such an increase could be quite large.

The results presented here offer solid evidence that for inputs in which the adult's time investment is a public good, a child's receipt of the input is increasing in the size of her age-cohort. Such findings suggest that there are economies of scale gained when a household's children are age-grouped rather

[10]The results in columns (3) and (4) may also pose the question, what is causing the increase in point estimates? However, the coefficients from the "uncorrelated" estimations are statistically indistinguishable from the original estimate at the 1% level.

than age-isolated. I offer the contribution of this work as a first look at scale economies arising from an aspect of household structure previously ignored by the economic literature. This paper presents evidence that such economies do exist.

References

A. P. Barten. *Econometric Analysis for National Economic Planning*, chapter Family Composition, Prices and Expenditure Patterns. Butterworth, 1964.

Angus Deaton and Christina Paxson. Economies of scale, household size, and the demand for food. *Journal of Political Economy*, 106(5):897–930, October 1998.

Earnst Engel. Die lebenskosten belgischer arbeiter-familien fruher und jetzt. *Bulletin de l'Institut International de Statistique*, (9):1–129, 1895.

Andrew D. Foster. Altruism, household coresidence and women's health investment in rural Bangladesh. Working Paper, Brown University Department of Economics, March 2004.

Li Gan and Victoria Vernon. Testing the barten model of economies of scale in household consumption: Toward resolving a paradox of deaton and paxson. *Journal of Political Economy*, 111(6):1361–1377, December 2003.

Peter Lanjouw and Martin Ravallion. Poverty and household size. *Economic Journal*, 105(433):1415–34, November 1995.

E. P. Lazear and R. T. Michael. Family size and the distribution of real per capita income. *American Economic Review*, (70):91–107, 1980.

John Muellbauer. Testing the barten model of household composition effects and the cost of children. *The Economic Journal*, 87(347):460–487, sep 1977. ISSN 0013-0133.

Robert A. Pollak and Terence J. Wales. Demographic variables in demand analysis. *Econometrica*, 49 (6):1533–1551, 1981.

S. J. Prais. The estimation of equivalent-adult scales from family budgets. *The Economic Journal*, 63 (252):791–810, 1953.

WHO. Vaccine preventable diseases monitoring system immunization schedules. Technical report, World Health Organization, December 2007.