

## ABSTRACT

Title of dissertation: CHILDHOOD EVENTS AND LONG-TERM CONSEQUENCES

Giordano Palloni, Doctor of Philosophy, 2015

Directed by: Professor Sebastian Galiani

Health and experiences in early childhood are strongly associated with adult outcomes. In this dissertation, I explore the association in detail with a focus on identifying the causal mechanisms that generate variation in early health and uncovering the parental behaviors that determine whether early health and living conditions evolve into long-term deficits or advantages.

In chapter 1, I explore whether pre-conception maternal desire for children of a particular sex has implications for the health of children in Indonesia. I show that a simple fertility stopping model predicts that when a child is born of the mother's preferred sex, they will receive more resources, and I test this prediction empirically using a longitudinal data set. I find that children born of the mother's preferred sex are heavier, have a higher body mass index, and experience fewer illnesses. I provide evidence that reductions in subsequent fertility are the primary mechanism for these effects.

The existing research measuring the long-term implications of early childhood conditions frequently fails to identify the mechanisms through which early deficits

become life-long disadvantages. In chapter 2, I examine one instance where deficits may matter for long-term well-being. Using data from Indonesia, I find that when third trimester rainfall is fifty percent higher than expected, birth weight and relative size are approximately .23 standard deviations higher. Despite this early advantage, I find no persistent positive impact fifteen years later. However, parental investment appears to be negatively influenced by in utero exposure to rainfall, suggesting that parents compensate for early health conditions.

To date, research on the long-term effects of childhood participation in subsidized housing has been limited by the lack of suitable identification strategies and appropriate data. In chapter 3, I, along with my co-authors, create a new, national-level longitudinal data set on housing assistance and labor market earnings to explore how children's housing affects their later employment and earnings. We find that while naïve estimates suggest there are substantial negative consequences to childhood participation in subsidized housing, household fixed-effects specifications attenuate these negative relationships for some demographic groups and uncover positive and significant effects for others.

EARLY CHILDHOOD EVENTS AND LONG-TERM CONSEQUENCES

By

Giordano Palloni

Dissertation submitted to the Faculty of the Graduate School of the  
University of Maryland, College Park in partial fulfillment  
of requirements for the degree of  
Doctor of Philosophy  
2015

Advisory Committee:

Professor Sebastian Galiani, Chair

Professor Jessica Goldberg

Professor John Haltiwanger

Professor Sangeetha Madhavan

Professor Sergio Urzúa

© Copyright by  
Giordano Palloni  
2015

## Foreword

The third chapter of my dissertation was co-authored with Fredrik Andersson, John Haltiwanger, Mark Kutzbach, Henry Pollakowski, and Daniel Weinberg and benefited enormously from their considerable ability and experience. This chapter was supported by grant number 98082 from the “How Housing Matters” research program of the John D. and Catherine T. MaCArthur Foundation and uses data from the Census Bureau’s Longitudinal Employer Household Dynamics Program, which was partially supported by: National Science Foundation Grants SES-9978093, SES-0339191, and ITR-0427889; National Institute on Aging Grant AG018854; and grants from the Alfred P. Sloan Foundation.

## Acknowledgements

This dissertation would not have been possible without the support of Sebastian Galiani, Jessica Goldberg, John Haltiwanger, and Sergio Urzúa. Their insight, patience, and willingness to engage me in serious discussions of preliminary research ideas helped mold early thoughts into a cohesive project. They deserve a great deal of credit for any novel or interesting material found herein.

Thank you to Sangeetha Madhavan for bringing a much needed demographer's perspective to my dissertation committee.

I owe a debt of gratitude to my co-authors: Fredrik Andersson, John Haltiwanger, Mark Kutzbach, Henry Pollakowski, and Daniel Weinberg. The third chapter of my dissertation benefited enormously from their considerable ability and experience. Thank you also to Emily Mytkowicz for research assistance and Daniel Hartley, Kristin McCue, and Erika McEntarfer for their comments on early drafts.

For guidance and human capital building exercises throughout graduate school, thank you to Raymond Guiteras.

To Rebecca Thornton, thank you for teaching me that good economic research is ultimately about interesting ideas.

A number of students and faculty members at the University of Maryland provided helpful comments and suggestions. Thank you to Brian Quistorff for detailed comments on all three chapters of this dissertation and on numerous other working projects. Thank

you to Lesley Turner, Melissa Kearney, John Ham, Miguel Sarzosa, and Tara Kaul for their time and notes during seminars.

For feedback on the prose and organization of early drafts, thank you to Samuel Hall.

Thank you to my parents for their constant guidance, love, and support.

And finally, thank you to Luisa, Sienna, and McCoy. You make everything worthwhile.

# Contents

<b>I Childhood Health and the Wantedness of Male and Female Children</b>	<b>1</b>
1 Introduction	1
2 Background	4
3 Conceptual Framework	10
4 Data	15
5 The Determinants of Maternal Gender Preference	22
6 Empirical Strategy	27
7 Results	29
8 Mechanisms: Realized Gender and Subsequent Fertility	38
9 Conclusion	45
10 Tables	48
11 Figures	58
12 Appendix A: Model	61
13 Appendix B: Additional Tables	69
14 Appendix C: Additional Figures	72
15 Appendix D: Identification of Treatment Effects	73
<b>II Parental Response to In Utero Shocks</b>	<b>78</b>
1 Introduction	78
2 Rainfall	83
3 Theoretical Framework	85
4 Data	92
5 Empirical Strategy	99



<b>6</b>	<b>Results</b>	<b>102</b>
<b>7</b>	<b>Conclusion</b>	<b>110</b>
<b>8</b>	<b>Tables</b>	<b>112</b>
<b>9</b>	<b>Figures</b>	<b>120</b>
<b>10</b>	<b>Appendix A: Model</b>	<b>124</b>
<b>III</b>	<b>Childhood Housing and Adult Earnings: A Between-Siblings Analysis of Housing Vouchers and Public Housing</b>	<b>127</b>
<b>1</b>	<b>Introduction</b>	<b>127</b>
<b>2</b>	<b>Literature Review</b>	<b>130</b>
<b>3</b>	<b>Research Design, Hypothesis, and Identification Issues</b>	<b>135</b>
<b>4</b>	<b>Description of the Data</b>	<b>139</b>
<b>5</b>	<b>The Sample: Basic Facts</b>	<b>145</b>
<b>6</b>	<b>Empirical Results</b>	<b>150</b>
<b>7</b>	<b>Extensions and Robustness Checks</b>	<b>155</b>
<b>8</b>	<b>Concluding Comments</b>	<b>157</b>
<b>9</b>	<b>Tables</b>	<b>160</b>
<b>10</b>	<b>Figures</b>	<b>170</b>
<b>11</b>	<b>Appendix A: Major U.S. Subsidized Rental Housing Programs</b>	<b>172</b>
	<b>References</b>	<b>174</b>

# List of Tables

## Chapter I

Table 1	48
Table 2	49
Table 3	50
Table 4	51
Table 5	52
Table 6	53
Table 7	54
Table 8	55
Table 9	56
Table 10	57
Table B1	69
Table B2	70
Table B3	71

## Chapter 2

Table 1A	112
Table 1B	113
Table 2	114
Table 3	115
Table 4	116
Table 5	117
Table 6	118
Table 7	119

## Chapter 3

Table 1	160
Table 2	161
Table 3	162
Table 4	163
Table 5	164
Table 6	165
Table 7	166
Table 8	167
Table 9	168
Table 10	169

# List of Figures

## **Chapter 1**

Figure 1	58
Figure 2	59
Figure 3	60
Figure C1	72

## **Chapter 2**

Figure 1	120
Figure 2A	121
Figure 2B	121
Figure 3	122
Figure 4	123

## **Chapter 3**

Figure 1	170
Figure 2	171

## Part I

# Childhood Health and the Wantedness of Male and Female Children

## 1 Introduction

The allocation of parental time and market resources across children has important economic and social consequences. Children who receive more investment in the form of nutrition, parental time, or educational inputs perform better across a wide range of outcomes during childhood and as an adult.<sup>1</sup> Though much of the variation in inputs and outcomes occurs across parental education and wealth categories [Currie, 2009], within household allocative inequalities are also highly predictive of a range of important differences across children [Jensen and Miller, 2011, Jayachandran and Pande, 2013]. As a result, explaining how and why parents allocate limited resources within the household is critical for understanding population-level distributions of health and well-being among adults.

When selecting optimal consumption and investment, parents may consider a number of different observable child characteristics. One dimension that has received substantial attention in the economics literature is how parental behavior varies with the sex<sup>2</sup> of children in the household. This research uncovers important male advantages in mortality, cognition, social and emotional skills, parental time, and anthropometric outcomes [Jayachandran and Kuziemko, 2011, Barcellos et al., 2012, Hu and Schlosser, 2011, Li and Wu, 2011], primarily through comparisons of population moments across females and males.

---

<sup>1</sup>For example, see Rosenzweig and Schultz [1982], Rosenzweig and Wolpin [1988], Case et al. [2005], Cunha and Heckman [2007], Qian [2008], Maluccio et al. [2009], Almond and Currie [2011], Jayachandran and Kuziemko [2011], Macours et al. [2012], Baker and Milligan [2013], Bharadwaj and Lakdawala [2013], Gertler et al. [2013]

<sup>2</sup>I use the terms gender and sex interchangeably throughout the paper.

These differences, notable in magnitude as well as in breadth, imply uneven treatment of male and female children in a number of countries. However, even in those societies where son-preference is widespread, a non-trivial fraction of parents will desire a daughter at some point in their family formation. As a result, comparing mean outcomes across sex will average over births where parents truly preferred to have son before conception and some where parents actually preferred to have a daughter. Although the resulting treatment effects are accurate measures of the population-level average differences, they ignore potentially interesting and important heterogeneity in the size of the effect. Further, such estimates do not capture the effect of being born of the parents' preferred sex on the time and resources they devote to each child. In fact, without the ability to identify parents' more- and less-preferred gender before the birth of each child, studies in countries with a no dominant sex-preference will wrongly conclude that child outcomes and the allocation of parental time and resources never depend on gender.

Using a panel data set from Indonesia, a country with no immutable son-preference, this paper investigates how being born of a mother's pre-conception preferred sex impacts the resources children receive in early childhood. Repeated observation of the same households over a fifteen year period enables the linkage of maternal preference for children of a particular gender to the realized gender of the next child born to that mother and that child's anthropometric (BMI-for-age and weight-for-age) and general health outcomes several years later. By conditioning on the observed, pre-birth gender preference of the mother, identification comes from comparing children born of the preferred gender of the mother to children in observably similar households with mothers who have the same pre-birth gender preference who happen to be born of the less-preferred gender. Under the assumption that gender-at-birth is not manipulated by parents, this is a natural experiment with a treatment probability of one-half.<sup>3</sup> To my knowledge, this paper is the first to match pre-conception maternal sex preferences to early childhood health outcomes to estimate the

---

<sup>3</sup>The average biological probability of having a boy is estimated to be 51.2%

impact of children being born of their mother's preferred sex.

The results indicate that pre-birth maternal preferences over child gender have economically and statistically significant effects on the resources allocated to children. Being born of a mother's preferred gender leads to an increase in body mass index (BMI), weight-for-age, and mother reported general health, and a decrease in the number of days the child is sick during the past month, in the incidence of wasting, and in the incidence of thinness. Thus, maternal preferences have real and important effects on measures of acute nutritional status later in childhood. I find no effects on measures that reflect longer-term nutritional deprivation (i.e. stunting). By combining the estimates with previous work [Victoria et al., 2008] linking weight and body-mass index at two years of age to schooling, income, and health in adulthood, the results suggest that there will be sizable long-term consequences to these acute nutritional deficiencies.

To explore what changes in parental behavior drive the main results, I link all mothers in the sample to their detailed fertility histories. I begin by showing that mothers who have a child of their preferred gender are significantly less likely to continue having further children and have significantly fewer subsequent children. Next, making use of a Two Sample Two Stage Least Squares (TS2SLS) strategy [Angrist and Krueger, 1992, Inoue and Solon, 2010], I combine estimates of the effect of the sex distribution of the first two children on family size with estimates of the effect of family size on BMI and weight-for-age. The results imply that reductions in subsequent fertility can explain roughly half of the total effect of being born of the mother's preferred gender on BMI and over one quarter of the effect for weight-for-age.

The remainder of the paper proceeds as follows. Section II provides background information on parental preferences over the sex of children and the discusses the importance of the anthropometric measures used as outcomes. Section III presents the conceptual framework and Section IV discusses the data and summary statistics. Section V explores the correlates of maternal preferences for children of a particular gender. Section VI and

Section VII present the empirical strategy and the results. Section VIII discusses potential mechanisms and Section IX concludes.

## **2 Background**

### **Background on Maternal Gender Preference**

To interpret the main results of this paper it is critical to understand what drives maternal gender preference. For a child to be born of their mother's more preferred (or less preferred) sex it must be true that their mother desired children of one sex more than the other before that child's birth. Therefore, the effects estimated in this paper could not exist if mothers were always indifferent about which sex they desire more for their next child. The literature on the determinants and consequences of parental sex preference dates back to the 1960s. Ben-Porath and Welch [1976] provides a theoretical basis for why parents might prefer children of one sex. This work builds on the discussion of household fertility and resource allocation in Becker [1960], Becker and Lewis [1973], and Becker and Tomes [1976]. All four papers focus on the fertility implications of parental sex preference in developed countries, though many of the mechanisms remain relevant for households in developing contexts. Jayachandran [2014] offers a more recent exposition, targeted specifically towards developing countries. In general, the potential causes of sex preference discussed in the literature can be assigned to one of three categories: economic, cultural and institutional, or individual taste-based.

One possible impetus for mothers to prefer that future children be of one sex is differences in the economic costs or returns to male and female children. For example, if parents believe that girls and boys require a different set of inputs, or that the optimal bundle of inputs differs by child gender, then the different market prices for these investment bundles are relevant. Intuitively, if high quality boys are more expensive to produce than high quality girls, even if parents value quality equally for both genders, they will prefer to have

children that are girls. Similarly, mothers may prefer to have male children if they expect to receive a higher return for male child labor or if they expect to receive more financial support from male children in old age [Rahman and Rao, 2004, Koolwal, 2007, Pande and Astone, 2007, Robitaille, 2013]. Brawn based economies,<sup>4</sup> economies where males have a comparative advantage over females, could contribute to mothers preferring male children over female children. Of course, if parents are able to reduce their investment in female children enough to equate the marginal return to resources across child gender they should be indifferent to the sex of subsequent children, instead leading to large gender gaps in health, human capital accumulation, and even mortality [Rosenzweig and Schultz, 1982, Qian, 2008, Pitt et al., 2012]. As this may require either active sex-selective abortion, infanticide, or massively unequal allocations of resources, it may not be feasible in countries where such practices are frowned upon or where parents are averse to inequality [Behrman et al., 1982, Pitt et al., 1990].

A second possibility is that cultural norms and institutions impact maternal preference for future sons and daughters. Numerous authors have suggested that the special role that sons take on in Hindu funeral rituals has contributed to the extreme son preference in India [Vlassoff, 1990, Pande and Astone, 2007, Jayachandran, 2014]. However, the gender gap in religious importance is smaller in Islam. As nearly ninety percent of Indonesians are Muslim and fewer than two percent are Hindu<sup>5</sup>, this suggests religion will be less important in the context of this study.

Often strongly related to religion, cultural norms associated with marriage might also contribute to maternal sex preference. For example, virilocality (patrilocality) or marital exogamy, traditions in which daughters move away from home upon marriage or daughters are married to individuals outside of their home village, have been proposed as potential mechanisms through which son preference might arise [Dyson and Moore, 1983, Yount, 2005]. However, recent work suggests exogamy may be of little empirical importance in

---

<sup>4</sup>Defined as economies with a high average economic return to physical strength.

<sup>5</sup>CIA [2014]



India [Rahman and Rao, 2004] and regional-level variation in virilocality is not responsible for differences in the treatment of sons and daughters in Indonesia [Levine and Kevane, 2003]. Dowries, or the price paid from the family of the bride to the family of the groom upon marriage, have also been identified as a possible factor influencing the relative demand for sons and daughters [Kishor, 1995, Arnold et al., 1998, Gupta et al., 2003]. As with virilocality, there appears to be little reason to believe that dowries influence the demand for or treatment of daughters in Indonesia. In fact, in contrast with many other South and East Asian countries, among some Indonesian ethnic groups there is a net transfer to the family of the bride at the time of marriage (bridewealth) [Errington, 1990, Boomgaard, 2003]. Similarly, customs that dictate which child bears the responsibility of caring for parents as they age could impact the relative demand for sons. In much of the developing world the eldest son is responsible for supporting parents as they age (see Dyson and Moore 1983, Ebenstein and Leung 2010). In the absence of social programs to provide care for the elderly, parents will desire at least one son to provide care. Again, Indonesia appears to be an outlier among developing countries with respect to the customs dictating responsibility for elderly parents: much of the population bequests the home to the youngest daughter in return for her caring for the parents in old age. This suggests, if anything, that there may be incentive for parents to prefer daughters over sons.

Societal rules related to property and land rights offer a third cultural/institutional channel which might lead mothers to desire that future children be of a particular gender. Both parents may prefer that land and possessions be passed on to their children upon their death. As such, religious or societal rules that deny or limit inheritance or bequests to female children could negatively impact the relative demand for daughters [Oldenburg, 1992, Dharmalingam, 1996, Carranza, 2012]. Prior to 1994 in Indonesia, there existed a conflict between civil and Islamic law with respect to gender-specific inheritance: civil law allowed for female children to be heirs and to exclude agnates<sup>6</sup> from inheritance while Islamic law

---

<sup>6</sup>Other relatives from the same male line.

allocated twice the female inheritance to all male heirs and prevented female children from excluding any other heirs [Lukito, 2006]. A 1994 ruling by the Indonesian government ended this inconsistency. In 1994, Islamic law was re-interpreted by the courts in Indonesia to also allow daughters to exclude other potential heirs. Further, the majority of households in Java, home to roughly 60% of Indonesia's total population, have never subscribed to the uneven inheritance rules prescribed by Islamic law, and instead, have typically bequeathed equal shares to male and female children [Errington, 1990, Brown, 2003]. This suggests that the vast majority of mothers included in this study should not desire future children to be male as a result of differential inheritance laws.

Finally, individual tastes not reflecting economic or cultural gender differences might contribute to maternal preference for children of a particular sex. Some mothers may simply derive more utility from daughters than sons or visa versa. Dahl and Moretti [2008] provide evidence from the United States consistent with this theory. Fathers are more likely to marry the mother of their first-born child, less likely to divorce conditional on marrying the mother of their first-born child, and more likely to have visitation rights conditional on being divorced if the child is a male. Survey and observed fertility responses support the hypothesis that son-biased utility is responsible for this gender gap in paternal participation. Of course, mothers need not have gender-biased utility functions to prefer that future children be of a particular sex at some point during their fertility. If the number of male children and the number of female children enter the mother's utility function as symmetric but separate arguments then under standard preference convexity assumptions mothers who have had more boys will desire more future girls and those that have had more girls will desire more future boys, all else constant. This desire for a balanced sex ratio among children is well established and has been used as an instrument for total fertility by a number of previous authors, most notably to estimate the effect of family size on parental labor supply [Angrist and Evans, 1998, Cruces and Galiani, 2007]. Consistent with parental desire for a balanced sex ratio among children, in all three countries studied (the United States,

Mexico, and Argentina), this research finds that parents who have at least two children are more likely to have a third child if the first two children were of the same sex. The few papers that directly estimate whether past gender realizations, typically parameterized as the male ratio among existing children, have any effect on stated parental desire for future male and female children use data from countries where sex-selective abortion, infanticide, and differential resource allocation have led to problems of “missing women” [Yount, 2005, Pande and Astone, 2007, Koolwal, 2007]. If, as seems likely, parental preferences are correlated with the likelihood of differential mortality then the estimated coefficient on the male ratio among existing children will be biased. However, in Indonesia, a setting with no contemporary gender gap in mortality [Kevane and Levine, 2003], I should be able to test whether maternal demand for sons and daughters is influenced by the male ratio among existing children.

While past research provides a useful outline as to what factors might influence maternal gender preference, ultimately, the cultural, institutional, and individual taste-based factors that impact the demand for male and female children among mothers in Indonesia is a testable empirical question. Section V uses data from Indonesia to check the relative importance of a number of different measurable characteristics including the male-ratio among existing children.

This paper identifies the effect of a child being born of their mother’s pre-conception preferred sex on early health outcomes. As I discuss in more detail later in the paper, each mother’s pre-conception preferred sex is identified using the self-reported desired number of future male and female children. While surveyors are instructed to interview subjects alone, for women to be eligible to participate in the survey module pertaining to desired future fertility, they must be currently married or have been married at some point in the past. Therefore, elicited preferences are potentially influenced by the preferences of the current or past husband(s). A simple empirical test confirms that there is a strong correlation between maternal and paternal preferences: in cases where both the mother

and father have a preference, that preference is the same 92% of the time. While this strong correlation could be the result of bargaining between spouses, it could also result from the matching of partners with similar tastes in the marriage market. Given the lack of available information on pre-marriage preferences, I elect to remain agnostic about the specific processes that result in the overwhelming agreement between maternal and paternal preferences over the gender of the next child. That said, all of the main estimates of the effect of being born of the mother's preferred sex are robust to also controlling for whether the child was born of the father's preferred sex.

### **Background on Childhood Body Mass Index and Weight**

To investigate whether being born of their mother's preferred sex affects the resources devoted to young children, I use two different categories of outcomes: anthropometric measures (BMI-for-age, Weight-for-age) and mother-reported measures of the incidence of different types of illnesses. Unfortunately, the data I use do not contain measures of parental or child time use or individual-level measures of consumption. Although the availability of these measures at the individual level would be ideal, inference based on them would be fraught with its own set of potential problems resulting from the likely presence of non-classical measurement error [Ahmed et al., 2005, Browning et al., 2014, Crossley and Winter, 2015]. I argue that BMI-for-age, weight-for-age, and measures of the incidence of illness accurately reflect the resources devoted to children and are predictive of other important long-term outcomes.

Whereas height is typically used as a measure of long-term nutritional deprivation, weight-for-height (BMI) and weight are generally thought to reflect more acute nutritional deficiencies [WHO, 2014]. Specifically, BMI and weight-for-age deficiencies (moderate malnutrition) result from one of three potential sources: inadequate caloric consumption over a period of days or weeks, the incidence of illness which lessens caloric absorption by the body, or a combination of the first two possibilities. Despite the fact that they only

reflect short- to medium-term deficits, weight-for-age and BMI are important measures of childhood well-being which are strongly associated with long-term outcomes. For example, there is evidence that a large fraction of childhood illness and illness-induced mortality in developing countries can be attributed to underlying moderate malnutrition [Pelletier et al., 1995, Fishman et al., 2004]. As such, it seems logical to interpret any observed differences in BMI and weight-for-age as manifestations of acute caloric deficits. Given the age of the children studied here (below the age of seven), their lack of decision making power in the household, and the likely balance in their unobservable characteristics including genetic differences resulting from the identification strategy I employ, the deficits represented by low BMI and weight-for-age can only be the result of differential resource allocation by parents.

Beyond their use as measures of the resources devoted to children and their association with childhood morbidity and mortality, BMI and weight-for-age in childhood are strongly correlated with important standards of adult well-being. Recent research links low weight-for-age and BMI in early childhood to lower levels of adult height, educational attainment, and economic productivity [Victoria et al., 2008]. Perhaps surprisingly, there is also a burgeoning literature linking childhood malnutrition (stunting, wasting, and underweight status) to chronic disease [Gluckman and Hanson, 2004] and heart disease [Barker et al., 1989, Eriksson et al., 2001] in adulthood.

### **3 Conceptual Framework**

To help illustrate why the pre-birth maternal preference for children of a particular sex might have real health consequences, Appendix A presents a simple formulation of the mother's choice problem. The intent of the model is to illustrate that even with no sex-specific prices of investment and no overarching maternal preference for children of a particular sex (with respect to either quantity or quality) there will be differences in child

health/quality between children that were born of their mother's preferred sex and those that were not. This section outlines the model assumptions and discusses the implications of the model which are shown in more detail in the appendix.

### **Modeling Mothers' Preferences**

I assume that mothers receive utility from the number of children they have  $n$ , the average health or quality of their children  $\bar{q} = \frac{1}{n} \sum_j q_j$ , and their distance from an ideal sex composition (male ratio) among their existing children  $d\left(\frac{m}{n} - \gamma\right)$ . Here  $m$  represents the number of male children and  $\gamma$  represents the mother's ideal male ratio  $\left(\frac{m}{n}\right)$ . In each period, mothers receive resources  $y$  which they allocate across all existing children through investment  $i$ . The price of investment is assumed to be equal for male and female children and is represented by  $p_i$ .

A direct result of the single price of investment, the constant marginal productivity of investment  $\left(\frac{\partial q_j}{\partial i_j} = 1 \forall j \in \{1, \dots, n\}\right)$ , and the equal marginal utility of quality across children  $\left(\frac{\partial u}{\partial q_j} = \frac{\partial u}{\partial q_k} \forall k, j\right)$  is that mothers will divide resources evenly across all existing children,  $\bar{q} = \bar{i} = i$ . In practice, there are likely to be some locality-level gender differences in the price of investment and some mother-level differences in the marginal utility mothers receive from investing in male and female children. However, the model assumes away these differences to highlight the potential effect of one particular mechanism through which a child's being born of their mother's preferred gender can lead to differences in resource allocation and health outcomes: the mother's fertility response to the gender of her most recent child. In addition, the model assumptions are consistent with the lack of evidence of systematic differences favoring male or female children found in the previous literature [Kevane and Levine, 2003] and in the empirical analysis undertaken later in this paper.

I assume that maternal utility is concave in the net benefit from an additional child. That is, it is concave in the sum of the direct utility increase from an additional child and

the indirect loss in average child quality resulting from the decrease in available resource per child. I also assume that the distance function mapping from the existing and desired male ratios to utility is increasing and concave.

Mothers, knowing the price of investment, the resources available for investment, the number of existing male and female children, and under the assumption that any future child will be male with probability  $\frac{1}{2}$ , state which gender they would prefer for their next child. They then observe the sex of their next child and elect whether to continue their fertility. The model makes predictions about how the existing number of male and female children influence stated maternal sex preference and how stated maternal sex preference and the sex of the next child (the child whose birth is the subject of my analysis) influence subsequent fertility and the health of the main child.

### **Stated Maternal Preferences Over Child Gender**

**Proposition 1:** Mothers will prefer that their next child be of the gender that takes them closer to a balanced sex ratio

In other words, mothers will prefer to have a daughter next if they have more existing sons than existing daughters and they will prefer to have a son next if they have more existing daughters than existing sons. This results directly from the assumption that  $\gamma$ , the ideal male ratio, is equal to one half. Note that this pattern is not consistent with several alternative theories. For example, if mothers systematically prefer children to be of the sex for which investment is cheaper, there should be no relationship between the existing male ratio and the desired sex of the next child as long as the existing male ratio is not manipulated by parents. Similarly, if inputs differ by sex and at least some inputs are partially public,<sup>7</sup> then it will be cheaper for parents to have children of whichever sex they already have more of. Thus, both of these alternative theories would suggest that parents prefer less balanced sex ratios.

---

<sup>7</sup>Sex-specific hand-me-down clothing would create one such input.

## **Continued Fertility and Resource Allocation**

**Proposition 2:** If a mother decides not to have a child in one period, then she will also elect not to have a child in subsequent periods.

**Proposition 3:** Mothers are at least as likely to continue fertility when they are farther away from a balanced sex ratio.

Proposition 2 follows directly from the fact that the random part of the model, the sex of the next child, is only relevant if the mother elects to have another child. As an implication of Proposition 2, there is no need to consider the fertility decisions of mothers who elected not to have a child in previous periods. Proposition 3 follows from concavity assumptions on the mother's utility function. That a mother would prefer to have a male child next does not necessarily imply that she will treat her next child better if it is a male. However, Proposition 3 suggests a mechanism which maps from stated maternal preferences and the realized gender of the next child to differences in the observed health of the next child. Further, this mechanism is present even without allowing for more preferred children to receive a larger share of household resources. This is because more preferred children will, on average, end up in smaller households.

**Proposition 4:** If the increase in marginal utility from an additional child is low, then mothers will elect not to continue fertility regardless of whether they have a balanced sex ratio.

**Proposition 5:** At lower birth parities, children who are born of their mother's preferred sex will be healthier (higher quality) than children who are born of their mother's less preferred sex.

Proposition 4 suggests that the mother's desire for a balanced sex ratio should only be relevant if the mother would actually consider having another child. If the marginal utility increase from an additional child, absent any consideration of the existing sex ratio, is



sufficiently low, then mothers will not have an additional child even if they are far from a balanced sex ratio. For example, in the IFLS data, mothers report their ideal total fertility. Birth parities beyond the mother's ideal number of children are likely to bring smaller marginal utility increases for the mother. Thus, Proposition 4 suggests that if the effect of being born of the mother's preferred sex works primarily through changes in fertility, the effects should be more pronounced at parities below the mother's ideal fertility level.

The model makes a clear prediction as to how an increase in fertility should affect the resources devoted to children: an increase in the number of children reduces the total resources available per child. This is essentially the quantity-quality tradeoff suggested in Becker and Lewis [1973, 1974], Becker and Tomes [1976].<sup>8</sup> In practice, as long as household members require non-zero resources in order to survive, an increase in the number of children alive in a household results in a reduction in the resources allocated to at least one family member.<sup>9</sup> In the context of Indonesian households having another child, this implies that either mothers must reduce their consumption or leisure [Angrist and Evans, 1998, Caceres-Delpiano, 2006] or they must reduce the resources or leisure allocated to one or more older children. If, as discussed in the Appendix, health is most sensitive to inputs at younger ages, then an equal reduction in resources for all children and parents would disproportionately affect the younger children in the household. Therefore, it seems reasonable that if all children experience a reduction in resources as a result of continued fertility, and children who are born of their mother's preferred sex at low birth parities are exposed to less later fertility, then these more preferred children will receive more resources

---

<sup>8</sup>This does not imply that household members will always be made worse off by an increase in household size. If at least one household member receives sufficient utility from the consumption or utility of other household members, it may be possible for all existing members to be made better off.

<sup>9</sup>If children have some expected productive value, then this need not be true. Given the age range studied (on average three and a half years of age), it is unlikely that these children have much current or past productive value. That said, parents may expect their children to have a productive value at some point in the future. If parents in Indonesia are able to borrow against the expected future productivity of young children, they may not need to reduce the resources allocated to any household member in the present. However, it seems unlikely that parents would be able to borrow enough to completely offset the increase in resources necessary to sustain the life of a new child, especially given the lengthy expected time period before the new child becomes a potentially productive asset.

and be healthier than those children that were born of the mother's less preferred sex, all else equal. This is exactly what Proposition 5 predicts.

Of course, even the empirical confirmation of these propositions does not rule out the possibility that there might be other reasons for children to be better off when they are born of the sex that their mother desired more. It might still be the case that in addition to changing their fertility behavior, mothers reduce their leisure or consumption or the leisure or consumption of older children in order to shift resources to the main child. In this paper, I am able to identify whether the continued fertility mechanism suggested by the model is present and then to quantify how much of the variation in outcomes can potentially be explained by this fertility mechanism.

## **4 Data**

### **The Indonesian Family Life Survey (IFLS)**

To estimate the effect of being born of the mother's preferred gender on child outcomes, I use RAND's Indonesian Family Life Survey (IFLS). The IFLS is a longitudinal household survey containing information on over 66,000 individuals. The first wave of the survey (IFLS1) was fielded in 1993 and subsequent waves were conducted in 1997 (IFLS2), 1998 (IFLS2+), 2000 (IFLS3) and 2007 (IFLS4). The IFLS is a representative sample drawn from thirteen provinces which contain 83% of the total population of Indonesia. Surveyors initially contacted 33,081 people across over 7,000 households. Subsequent waves were conducted in . In later waves, attempts were made to contact all of the initial households as well as any new household members. In all cases, surveyors attempted to interview respondents alone. The IFLS contains detailed information about household consumption and the educational attainment, work history, marital history, and fertility history of all adult household members. In addition, detailed health measurements are available for most household

members. All of the health measurements, including height, weight, and hemoglobin level were obtained by trained nurses. Each wave also contains surveyor and mother reported general health, an illness history, and inpatient and outpatient medical center utilization for each child.

Crucially for this analysis, each wave of the IFLS asks married adult subjects about all existing and desired children. Data on existing children are collected through a fertility history reported by the wife. By aggregating over all four survey waves, I am able to build a complete fertility history for each mother including the sex, date of birth, and place of birth for each child. Husbands and wives are queried separately about their desired future fertility. Specifically, subjects are asked the following questions pertaining to desired fertility: “How many (more) children do you wish to have?” and “Among the children that you (still) wish to have, how many sons and daughters do you wish to have?” I classify mothers as having a son-preference if they desire more future sons than future daughters and a daughter preference if the reverse is true. Mothers who desire no future children or equal numbers of each gender are classified as having no gender preference.

The panel structure of the IFLS data enables me to link elicited maternal preferences to information about the next birth for that mother. As a result, I observe whether the realized gender of the next child born following the measurement of maternal preferences matches the stated preference of the mother. I construct an indicator variable equal to one if the mother reported preferring future female children and her next child was born female or if the mother reported preferring to have future male children and her next child was born male. This indicator is equal to zero in all cases where the mother’s expressed preference does not match the realized gender. This measure of whether children were born of the sex that their mother preferred before their conception is the explanatory variable of interest. It is important to note that this measure is distinct from the effect of being born to a mother with a particular gender preference.

All anthropometric outcomes are standardized using the 2006 World Health Organiza-

tion (WHO) Child Growth Standards [WHO, 2006] for children under five years of age and the corresponding standards for children aged five to eighteen years [de Onis et al., 2007]. The WHO constructed the 2006 measurement and growth standards (CGS) by following 8,440 healthy, breastfed children of non-smoking mothers from Brazil, Ghana, India, Norway, Oman, and the United States. The CGS are intended to represent the growth patterns of children in ideal circumstances. The 2007 measurement growth standards (CGS2007) are constructed by smoothing data from the 2006 WHO growth standards for children under five and a National Center for Health Statistics/WHO growth reference for individuals under age twenty-four. Combining the two sets of growth standards yields monthly standards for boys and girls from birth to eighteen years of age.<sup>10</sup>

For each gender and age-in-month cell, the standards provide a median measurement, a coefficient of variation, and a measure of Box-Cox power for height, weight, and BMI. I construct Z-scores following the suggested WHO methodology by merging measurements from children in the IFLS by month of birth and gender to the CGS.<sup>11</sup> As a result of the standardization, effects should be interpreted as standard deviations of the reference distribution.

As with previous research employing anthropometric measures as outcomes, one potential caveat is that the setting-specific distribution of anthropometric measures might differ by gender.<sup>12</sup> For example, Indonesian women may be taller and heavier when compared to international standards for women than Indonesian men are when compared to international standards for men. To account for this possibility, I include a male indicator in all specifications with controls. In practice, the results are identical to those in the main specification with no controls even when including a male indicator fully interacted with

---

<sup>10</sup>There are no children in the sample over the age of eight.

<sup>11</sup>To limit the influence of outliers, Z-scores are top (bottom) coded to 5 (-5). The results are not statistically different when I do not top and bottom code that data. Similarly, dropping observations that would have been top or bottom coded has no effect on the results.

<sup>12</sup>In my sample, the standardized outcomes do not have significantly different distributions for male and female children. A Kolmogorov-Smirnov test fails to reject the null hypothesis that there is no difference in the distribution of BMI-for-age and weight-for-age by sex. Results available upon request.

a quadratic polynomial in age-in-months. Further, results are also identical in sign and significance when using similarly constructed deviation measures that employ only within-sample variation from each gender by age-in-month cell.<sup>13</sup>

Additional outcome variables include mother reported measures of the general health of each child, and two mother-reported measures of illness for each child: how many days of their primary activity the child missed and how many days the child was bedridden during the four weeks preceding the survey. Finally, I construct indicator variables that measure whether the child is wasted, thin, or stunted. All three measures are also defined according to the WHO growth standards: a Z-score below -2 for height-for-weight (BMI), weight-for-age, and height-for-age, respectively.

In some specifications, an additional set of control variables is included. These variables include an indicator for whether the child is male, indicators for whether the mother preferred to have a boy or a girl, age-in-months, year of birth indicators, month of birth indicators, province indicators, a linear trend which is allowed to differ by province, an indicator for whether the household lives in an urban setting, the mother's age, survey wave indicators, total household size including adults, and a full set of fixed effects for the number of older brothers and the number of older sisters in the household.

I use the fertility history for each mother to identify whether the sex of the first two children born to mothers is the same (i.e. both were girls or both were boys). In addition, I observe whether they had children after the child included in the main sample. I create an indicator variable for whether the mother had at least one additional child and another variable with the total number of additional children birthed by the mother. For those mothers who continue their fertility, I also generate a variable containing the time, in months, between the birth of the main child and the birth of the subsequent child.

For ease of exposition, in the remainder of the paper children who are the more preferred gender of their mother are referred to as MP; Children who are the less preferred

---

<sup>13</sup>The point estimates are different in magnitude as they represent deviations from a different distribution. That said, the estimates are identical in sign and statistical significance.

gender of their mother are referred to as LP.

### **Summary Statistics**

Panel A of Table 1 displays summary statistics for the main analysis sample. To be included in the main sample, children must be the first born to their mother after she responds to the future fertility preference questions discussed above.

Just over 22% of children were born to mothers who would prefer to have a female child and another 22% were born to mothers who would prefer to have a male child. The remaining 56% of children were born to mothers with no stated preference for children of either gender. Combining existing and desired future children, mothers desire roughly 1.039 male children for each female child in total. In 2004 in India, the corresponding ratio was 1.39 male children for each female child. Thus, there is significantly less son-preference in Indonesia than in India. On average, mothers desire 1.2 more children than they already have. The mean number of older brothers and older sisters for the children in the main sample is .93 and .88, respectively. 50.7% of the children in the sample are male, which is not statistically significantly different from the estimated biological likelihood of having a male child: 51.2%.

Mothers have completed 8.2 years of schooling on average. During the week before the survey, just under 60% of mothers were primarily involved in housekeeping while another 39.1% spent most of their time working. Over 90% of mothers are Muslim, another 5% are Christian and nearly all of the remaining 5% are Hindi.

On average, children were too sick to participate in normal activities 1.67 days and bedridden .42 days out of the past twenty-eight. As expected, children are considerably more likely to be wasted (BMI), lighter (weight-for-age), and shorter (height-for-age) than children in the reference sample. Defined as having a Z-score below negative two, 11% of children in the sample are classified as wasted (BMI), 26 percent are classified as thin (weight-for-age), and nearly 39 percent are classified as stunted (height-for-age).

Panel B of Table 1 displays summary statistics for the limited variables available for children appearing in a mother's fertility history. To be included in the sample, the requirements are minimal: children must have a non-missing year and month of birth, sex, and a mother who completed a fertility history at some point. The purpose of this sample is to test for sex-biased behavior on the part of mothers and to help estimate the effect of family size on health outcomes later in the paper.

Table 2 compares a subset of the observable characteristics across three different classifications of maternal gender preference: mothers with no gender preference (NGP mothers), mothers who prefer sons (SP mothers) and mothers who prefer daughters (DP mothers). Column 1 contains the means for the children of NGP mothers, Column 2 for SP mothers, and Column 3 for DP mothers. Columns 4 and 5 show the differences and p-values from tests of whether there is no difference between columns 1 and 2. Columns 6 and 7 show the same differences and p-values for the comparison between NGP children and DP children, and Columns 8 and 9 do the same for the comparison between SP children and DP children.

There are a number interesting differences across the three groups. First, maternal gender preferences are highly negatively correlated with the gender ratio of previous children. SP children have, on average, only .26 older brothers. The corresponding averages for NGP children and DP children are 1.043 and 1.343. The means for NGP and DP children correspond to 4.01 and 5.16 times the mean for SP children. Similarly, DP children have significantly fewer older sisters than SP children and NGP children. Converting the older sibling counts into the ratio of male to female ratios, the average for NGP children is 1.063, the average for SP children is .208, and the average for DP children is 5.415. These differences in means do not represent causal effects, but the strength of the negative correlation between the son ratio of existing children and the demand for further sons is quite striking.

Comparing the likelihood that children of each maternal sex-preference type are male reveals one important difference: SP children are significantly less likely to be male than

children from the other two groups. Although biological birth patterns dictate that the fraction of children that are male be .512 for each of the three preference types, this difference is not terribly troubling for my identification strategy. Specifically, these differences are not consistent with the primary concern that parents are manipulating the realized sex of their subsequent child. If subsequent child gender were endogenous to maternal preferences, we would expect the relationship to be the reverse. That is, SP children should have a higher likelihood of being male. Despite this, later in the paper I check whether the differences might be due to differential mortality.

NGP children have mothers who are younger and less educated than both SP and DP children. There are no significant differences with respect to maternal years of education completed or age between the SP and DP groups. The mothers of SP children are less likely to be Muslim and more likely to be Hindu. This is consistent with evidence from Northern India presented by Bhat and Zavier [2003].

Despite the observed differences in maternal education, age, and household composition, there is only one significant health difference across the three categories. DP children are, on average, reported as being less healthy than the children of SP and NGP mothers. The fact that there are differences in observable characteristics across NGP, SP, and DP mothers does not pose problems for my identification strategy. Instead, it illustrates why testing for the direct effect of different maternal sex-preference categories on child health outcomes would be problematic: there is important selection into each of the three categories.

Finally, Table 3 presents the raw, difference-in-means analog to the main regression results presented later in the paper. Limiting the sample to GP children, Column 1 displays means for LP children and Column 2 displays the corresponding mean for MP children. I compare 52 different observable household and characteristics but to ensure the results are legible Table 3 displays only 17. The remaining 35 means, differences, and P-values are available on request. Of the 52 non-outcome variables tested, in only one case are



the means significantly different from each other. LP children are slightly more likely to be born in March (not shown). This difference is significant at the 10% level. The other 51 comparisons across the two groups are not statistically significant. This is strong evidence that parents are not manipulating the sex of children through either active (abortion or infanticide) or passive disparate treatment. Further, it suggests that MP and LP children are well balanced on observable baseline characteristics. Thus, after matching maternal preferences to subsequent child gender, Table 3 suggests I can proceed as though treatment (whether a child is born of their mother's preferred sex) is randomly assigned with a treatment probability of one-half.

In a preview of the main estimation results, the bottom panel of Table 3 shows there are a number of statistically significant raw differences in outcomes across MP and LP children. MP children missed fewer days of activity due to illness over the previous four weeks, they have higher BMI Z-scores, higher weight-for-age Z-scores, and are less likely to be wasted. Most of the remaining differences in outcomes are marginally insignificant but signed as to suggest that MP children are healthier than their LP counterparts.

## **5 The Determinants of Maternal Gender Preference**

To build on the discussion in Section IIA, this section undertakes an empirical analysis of the determinants of stated maternal sex preference for the mothers in my sample. In particular, I focus on testing the relative importance of the sex distribution among existing children as compared to other individual, economic, cultural, and institutional factors. In addition to providing an interesting insight into the nature of sex preference, an understanding of why mothers desire that future children be male or female is critical for any policy implications. As noted above, previous empirical papers have been unable to estimate the causal impact of the existing sex ratio on stated sex preference. As a result, the prescribed policies focus on altering the individual, institutional, and cultural characteristics that influence sex

preference. For example, if the relative wages of women are important determinants of the demand for daughters, then policies that improve this relative wage should also increase the demand for daughters. Alternatively, if gender-biased property or inheritance laws that prevent daughters from owning historically familial lands or goods explain son preference, then legislation that alters these laws could potentially reduce stated son preference. On the other hand, if the sex ratio of existing children is the primary driver of the desire for future children to be of a particular sex, it is less clear what policies could be implemented to reduce any resulting disparities in resource allocation. If mothers are unaware of the real consequences of the differential allocations, providing sound estimates of these parameters could potentially lead them to reduce the disparity in treatment. Alternatively, policies or transfers that directly provide health and human capital inputs to children during critical periods of their development could reduce the negative consequences resulting from any within-household allocative disparities that arise as a result of the maternal preference over the sex of future children.

To interpret the coefficient estimates on the existing sex ratio as causal, it must be the case that the existing sex ratio is uncorrelated with any unobserved heterogeneity. Although ultimately untestable, at a minimum this requires that there is no sex-selective abortion, infanticide, or mortality resulting from either conscious or passive unequal resource allocation. Any version of sex-biased early mortality would skew the observed existing sex ratio. Further, the skewed sex ratios would occur in households with the strongest sex preference, and therefore would bias estimates. To begin, I calculate the fraction of children who are male at each birth order between one and seven.<sup>14</sup> Figure 1 plots this fraction by the birth order of the child as well as a 95% confidence interval for each parity. In addition, Figure 1 displays a horizontal line at the biologically expected fraction, .512. This corresponds to a ratio of 105 male children for every 100 female children. At no parity between one and seven does the fraction of children who are male differ from the biologically expected frac-

---

<sup>14</sup>Birth orders above seven represent the top one percent of the distribution. Including these observations does not affect the results. I elect to omit them for ease of exposition.

tion. This suggests there is little or no sex selective mortality in the sample and, therefore, that the sex ratio among previous children can be included as an exogenous explanatory variable.

Panel A of Table 4 turns to exploring the correlates of stated maternal sex preference. I construct a measure of future son preference by taking the ratio of the number of desired future sons to the total number of desired children. This ratio is bounded between zero and one with values above one-half indicating the mother would prefer to have more future boys and values below one-half indicating the mother prefers to have more future girls. I also construct the analogous measure for existing children by dividing the number of existing male children by the total number of existing children.

Next, within each of the thirteen provinces included in the data, I generate a measure of the average value of the bride wealth provided to the bride's family at the most recent marriage. I also calculate the province level fraction of women who moved away from their village of residence immediately following their last wedding. Together, I interpret these variables as a summary measure of the expected value of daughters in a given geographic area.

To provide insight as to what influences whether a mother prefers a male or a female child next, I regress the measure of maternal son preference on the fraction of existing children that are male, the province-level measures of average bride wealth, and the likelihood of moving post-wedding. In addition, I include a full set of number-of-children dummies, a linear time trend, month of birth indicators, dummies for whether the mother is Muslim, Hindu, or Christian, a set of indicators for the mother's highest education level, a dummy for whether the mother worked last week for pay, a quadratic in the mother's age, a continuous measure of total household size, a continuous measure of the household's expenditure and production of food per capita (in thousands of Rupiah), the total earnings of the mother during the past year (in hundreds of thousands of Rupiah), and two measures of maternal health: body mass index and height in meters. Standard errors are clustered

at the kabupaten (district) level. I emphasize that interpreting the regression coefficients other than the coefficient on the existing male fraction as causal parameters would require making strong assumptions about unobserved heterogeneity and the direction of the relationship between each control and the dependent variable. That is not the intent of this exposition. Instead, the aim is to highlight the strength of the partial correlation between maternal gender preference and different cultural, religious, and household level factors.

Column 1 displays the estimated coefficient on the fraction of existing children that are male, and the R-squared from the regression after controlling for all the characteristics discussed above. The observed characteristics are highly predictive of maternal gender preference with an R-squared of over .6.

Nearly all of the estimated coefficients have intuitive signs. For example, a higher average bride wealth and a lower likelihood of the bride moving to a new village immediately following the wedding are associated with a decrease in stated son preference. Similarly, mothers may desire more daughters if they expect to remain in closer geographic proximity to their daughters after marriage. The negative time trend suggests that son-preference has decreased over the study period.

Muslim and Christian women display lower son preference than Hindu women. This is consistent with data from other Southeast Asian countries and potentially attributable to the belief among Hindus that ascension to heaven requires a son to light a funeral pyre. Higher maternal education and household food expenditures per capita are associated, albeit weakly, with stronger son preference. Conversely, there is a significant negative correlation between the labor earnings of the mother during the past year and her son preference suggesting that the economic opportunities available to women may play a role in determining the demand for female children. A one hundred USD (1.14 million Rupiah) increase in the mother's salary last year is associated with a statistically significant decrease of 1.2% of the mean desired son ratio for the sample.

Both height in meters and BMI are positively associated with son preference in the

sample. This suggests that mothers may not infer that their own health is indicative of their genetic advantage at producing high quality female children. Instead, these health indicators may be picking up residual effects of household wealth not captured by the similarly signed coefficients on maternal education and weekly household food expenditure per capita.

Column 2 displays the estimated coefficient on the fraction of existing children that are male, the chi-squared statistic and p-value from a test of whether the coefficient on the existing son ratio is equal with and without controls, and the R-squared from the regression of the desired son ratio on the existing son ratio with no controls. Though several of the variables included in Column 1 have a statistically significant relationship with the desired son ratio, nearly all of the variation is explained by the fraction of existing children that are male. The R-squared without controls is .589, just .022 below the R-squared with the full set of controls. Consistent with mothers having no strong son or daughter preference, mothers who have a higher number of existing sons, holding constant the existing number of daughters, desire fewer future sons. The point estimate suggests that changing the sex of one existing child from female to male reduces the desired number of future male children by nearly 3/4 of a child. This relationship is significant at the 1% level. Further, the coefficient on the existing son ratio is not affected by the inclusion of the extensive set of controls. A test of equality of the effect of the existing son ratio across the two specifications fails to reject the null hypothesis of no difference (p-value of .855). The coefficient stability when including controls provides additional evidence that the estimated coefficient on the existing son ratio is capturing a causal parameter.

To further illustrate the explanatory power of the existing son ratio, Panel B of Table 4 uses a linear probability model (LPM) to predict maternal son preference in the main sample, with and without controls. Predicted values from each LPM are generated and mothers are classified on the basis of these predictions. Mothers with a predicted probability above one half are classified as being likely to prefer a boy while those with a predicted proba-

bility below one half are classified as unlikely to prefer a boy. These predictions are then compared to the observed maternal son preference to gauge the accuracy of the LPM. The LPM controlling only for the existing son ratio correctly predicts son preference 83.4% of the time. Adding the full set of controls barely improves upon this figure, correctly predicting son preference 84.5% of the time. A t-test of whether the two specifications are equally accurate fails to reject the null hypothesis of no difference (p-value of .289).

Panels A and B of Table 4 and emphasize that the gender distribution among existing children is the most important determinant of maternal gender preference for mothers in the sample. Also, despite there being no immutable gender preference in the sample, at any given point in time a large fraction of mothers do prefer to have their next child be born of a particular gender. This preference is largely dictated by the gender distribution among existing children.

## 6 Empirical Strategy

Given the evidence presented above that a child’s sex at birth is unrelated to family characteristics and pre-birth maternal sex preferences, the main analysis of the paper assumes treatment is as good as randomly assigned. A proof mapping from the average treatment effect estimator and identification assumptions to the empirical implementation described here is available upon request. I estimate the following model for child  $i$  born to mother  $m$ :

$$y_{im} = \alpha + \beta Pref_{im} + \eta X_{im} + \varepsilon_{im} \quad (1)$$

The outcome of interest  $y_{im}$ , depends on a constant an indicator for whether the child  $i$  was born of their mother’s preferred gender ( $Pref_{im}$ ), and a vector of controls ( $X_{im}$ ). In my preferred specification,  $X_{im}$  is empty and (1) is simply the Ordinary Least Squares regression of the outcome on the treatment indicator and a constant. In specifications with controls,  $X_{im}$  includes an indicator for whether mother  $m$  preferred to have a female child

immediately prior to the birth of child  $i$ , a male dummy, the child's age in months, the mother's age, a full set of dummies for the number of older brothers child  $i$  has, a full set of dummies for the number of older sisters child  $i$  has, household size including adults, year of birth dummies, month of birth dummies, province dummies, a linear trend which is allowed to differ by province, an indicator for whether the household lives in an urban area, and a set of indicators for the wave of the survey the observation comes from. The identifying assumption is that whether the child is born of their mother's preferred gender is independent of  $\varepsilon_{im}$ . As discussed in the introduction, if the sex at birth of child  $i$  is random then this assumption will be satisfied. Further, because the basic specification includes no controls despite the extensive set of child, mother, and household characteristics available in the data, I am able to test whether the estimated treatment effects are affected by the inclusion of the full controls. It is extremely unlikely that there could be unobserved heterogeneity that is correlated with both the treatment indicator and the outcome but not with the control variables. I first estimate the preferred specification with no controls. I next estimate (1) including the full set of controls in  $X_{im}$  and test whether the estimated  $\beta$ 's are equal. If the identifying assumption is correct, including a detailed vector of further controls should not affect  $\beta$ , though the precision of the estimates is likely to change. To maintain a consistent sample across the specifications with and without controls, I limit the sample when estimating (1) to children with non-missing controls. In practice, few children are missing controls so this restriction has no impact on the coefficient estimates or their precision.

Because the data are pooled across all three waves of the IFLS, some mothers contribute multiple observations. Therefore, standard errors are clustered at the mother level. Ideally, the data would contain multiple observations from all mothers and mother fixed effects models could be estimated. I have experimented with running mother fixed effects models but there is not sufficient within-mother variation to generate precise estimates. However, the point estimates from the within-mother specifications are similar in magnitude and sign

to the less restrictive OLS estimates that use both between- and within-mother variation.

In practice, equation (1) could be estimated in three different ways. The sample could be limited to the children of mothers who preferred one gender and the preference indicator or the constant could be omitted. Alternatively, (1) could be estimated using all children matched to maternal preferences, with both preference indicators, a constant, a male dummy and with  $Pref_{im}$  defined such that NGP children of NGP are treated as MP children. Finally, (1) could be estimated using all children, including a constant, both preference indicators, a male indicator and defining  $Pref_{im}$  such that NGP children are treated as LP children. In practice, all three methods produce mechanically identical estimates of  $\beta$ . This is because after including the set of maternal preference indicators, there is no variation in the  $Pref_{im}$  variable within the group of NGP children. Thus, the variation for this group is completely absorbed by the preference indicators and these children do not directly contribute to the identification of  $\beta$ . However, if the effect of any of the control variables that are included is different for NGP children than for SP and DP children, estimates of  $\beta$  could be slightly different for these two samples. As the main specification does not include additional controls this becomes less of a concern. Ultimately, I elect to exclude the NGP children to make the interpretation of the estimates as simple as possible.

I test for effects with a number of outcome variables intended to capture short, medium, and long term differences in resource allocation. Specifically, I check whether children who are born of their mother's pre-conception preferred sex have higher BMIs, are heavier, taller, and less sickly using the variables described in Section IV.

## 7 Results

### Main Results

Table 5 displays the results from estimating (1) with and without controls. As explained above, each column displays the constant and the coefficient on  $Pref_{im}$  from an OLS re-



gression. The top panel specifications have no additional control variables while the bottom panel shows results from specifications that include the full set of controls. The outcome variable is listed at the top of each column. Finally, the p-value from a test of whether the coefficient on *Pref<sub>im</sub>* is equal with and without controls is presented at the base of each column.

MP children have higher BMI Z-scores than their LP counterparts. On average, MP children have BMIs that are .236 standard deviations higher, a difference that is significant at the 1% level. Adding the full set of controls has no effect on the estimated treatment effect (p-value .913). A .236 standard deviation increase is roughly equivalent to the mean difference in BMI-for-age Z-score when comparing the children of mothers with no formal education to the children of mothers with at least some middle school education (.243 standard deviations). Not surprisingly given the close relationship between BMI and weight, column 2 indicates that weight-for-age is also significantly higher among MP children. Being an MP child leads to a .164 standard deviation increase in weight-for-age. Further, column 6 shows that this variation is largely driven by increases in BMI among children at the bottom of the BMI distribution: MP children are 3.2 percentage points less likely to be wasted (Z-score below -2). Only 10.3% of the estimation sample is wasted implying that being MP implies a 31% reduction in the likelihood of being wasted at the sample mean. Together, the weight-for-age, BMI, and wasting results suggest that MP children are significantly bigger on average and that this difference is driven by reductions in the likelihood of moderate malnutrition.<sup>15</sup>

Consistent with the anthropometric evidence, mothers report MP children are healthier along a number of different metrics. On average, the mother reported general health of the child is nearly a tenth of a standard deviation lower (healthier) for MP children than LP children. Similarly, mothers reported that MP children missed nearly half a day less of their primary activity and were bedridden for 0.16 fewer days over the four weeks preceding the

---

<sup>15</sup>Moderate malnutrition implies a child is either stunted, wasted, or both.

survey. These differences are significant at the 1% and 10% levels, respectively. Finally, MP children were significantly less likely to have experienced a cough or a fever over the same period.

The regression results support the conjecture that MP children receive more household resources in early childhood than LP children. Importantly, they are less likely to be moderately malnourished,<sup>16</sup> a condition which, in addition to having direct consequences for the productivity and skill accumulation of children, is responsible for approximately 50% of the childhood mortality in developing countries [Pelletier et al., 1995, Caulfield et al., 2004]. MP children are significantly bigger and otherwise healthier than LP children.

The inclusion of controls does not significantly change the estimate of the treatment effect for any of the eight outcome variables. Together with the lack of observable differences, this supports the idea that whether a child is born of their mother's pre-conception preferred sex is exogenous to potential child health outcomes.

### **Robustness and Heterogeneity**

The main results point to an important relationship between pre-conception maternal sex preference and health in early childhood. Below I investigate how robust the main results are to a number of additional restrictions. In addition, I test for heterogeneity along observable dimensions with the intention of ruling out alternative explanations and identifying potential mechanisms.

### **Selective Mortality**

Section V shows that there is no sex-biased mortality in the sample. However, because the sample is split evenly between mothers who preferred to have a son and mothers who preferred to have a daughter, differential mortality between MP and LP children is still

---

<sup>16</sup>I elect not to show the height-for-age results because of space constraints. In neither panel does the estimated coefficient on  $Pref_{im}$  suggest any relationship between height-for-age and being born of the mother's pre-birth preferred sex.

potentially problematic. Both prenatal and post-birth differences in treatment could in theory generate an MP/LP mortality gap as prenatal sex-diagnostic technology was certainly available for at least some households in the sample.

To understand why differential mortality would be problematic, consider the case where prenatal mortality is artificially higher for LP children than MP children. If this differential mortality occurs primarily in households with high potential child outcomes, as might be the case if only “better” households had access to sex-diagnostic technology, there would be more MP than LP children who survive to birth. Further, this difference would be driven by children who had high potential outcomes. This would cause the main results to wrongly attribute this positive selection for the MP group to the treatment effect leading to a positively biased estimate. The same logic applies in the event of differential post-birth mortality.

Therefore, to assess whether selective mortality is a problem, I test for both prenatal and post-birth differences in mortality across MP and LP children. I use maternal reports of miscarriages, abortions, non-live births, and live births that did not survive to the date of the survey. From these reports, I generate two indicator variables: 1) whether the mother had a conception that did not survive to term for any reason and 2) whether the mother conceived a child that did not survive to the date of the survey. These measures, which both cover the time period between when the mother stated her sex preferences and the next survey wave, are intended to capture prenatal selection and any type of differential selection, respectively. If differential mortality is important, I would expect to estimate a positive, significant relationship between the likelihood of having a child of the mother’s preferred sex and at least one of the two early mortality measures.

Columns 1 and 2 of Table 6 present these results. All columns present OLS estimates from a bivariate regression of an outcome on the treatment variable and a constant, the same specification used in the “Without Controls” panel of Table 5. In both columns 1 and 2, the estimated coefficient is not significantly different from zero and extremely small

in magnitude. This suggests differential mortality is not confounding the main results. Further evidence of this point was presented in Table 2. As discussed above, any differential mortality should lead to there being more MP children than LP children. In fact, there are fewer MP children in the sample. Thus, the available evidence suggests that differential mortality does not pose a problem for the main estimates.

### **Early Conditions**

With access to sex-diagnostic technology, mothers could begin differential treatment of children while still pregnant [Bharadwaj and Lakdawala, 2013]. Even if the disparity is not severe enough to cause a mortality gap, it could create a gap in development that would eventually lead to starkly different health outcomes by *Pref<sub>im</sub>* group. However, as long as there is no mortality gap, differential prenatal treatment would not bias estimates. Instead, it provides one plausible mechanism through which pre-conception maternal sex preference could impact health in early childhood.

Fortunately, the IFLS contains information on both the birth weight and gestational length of most children. There is a well-established relationship between birth weight, gestational length, and later-in-life outcomes including health [Barker and Clark, 1997, Black et al., 2007, Almond et al., 2009, Almond and Currie, 2011]. Given the similarly well-known link between maternal nutrition and birth outcomes [Wu et al., 2004, Abu-Saad and Fraser, 2010, Milazzo, 2014], there is good reason to believe differences in maternal nutrition would lead to differences in early childhood health. To test whether prenatal differences in treatment might help explain the results, I estimate (1) using birth weight and gestational length as outcomes.

Columns 3 and 4 of Table 6 present the results. If differential treatment occurs in the womb, I would expect to see positive, significant estimates of the effect of being born of the mother's preferred sex. If, however, mothers do not observe the sex of children before they are born or maternal nutrition is independent of fetal sex, there should be no signifi-

cant relationship. Table 6 strongly supports the idea that there are no prenatal differences in treatment. In neither column is the estimated effect significantly different from zero. Further, the effect of being born of the mother's preferred sex on birth weight is actually negatively signed, the opposite of what would be expected. In sum, it does not appear that the main effects are the result of differential investment and resource allocation before birth.

### **Misclassification of Mothers**

Another potential concern is related to the method used to categorize mothers by their gender preferences. Specifically, it is possible that some mothers may be incorrectly assigned. Mothers who desire at least one more child of each gender, but more boys than girls, will be classified as having a preference for boys. It may be the case, however, that the mother would actually prefer to have the girl first, followed by the boys. Because subjects are not asked what gender they would prefer for their next child, there is no way to avoid this type of misclassification with this data set. That said, despite being non-classical, this type of measurement error should bias my estimates against finding any significant differences. Consider the case of a mother who desires two future boys and one future girl who prefers to have the girl first. For simplicity, also assume that the treatment effect is homogeneous. Clearly, if the true treatment effect is zero, this misclassification will not impact the estimates: they will still be zero. If the true treatment effect is positive, so that children who are born of their mother's preferred sex end up receiving more resources and therefore end up healthier, this assignment error will bias the estimates downwards by missclassifying some LP and less healthy children as MP children. Thus, the impact of the misclassification is to bias my results against finding any effect.

The data also permit me to test empirically whether this type of misclassification is important. In practice, nearly all mothers who I classify as having a preference for one gender desire only children of that gender. Of the 1518 cases where mothers are classified

as preferring either girls or boys, over 86% desire only girls or only boys. To test whether including potentially misclassified mothers is important, Table 7 displays the results when limiting the sample to cases where the mother desired that future children be of one gender only. As expected given the paucity of potentially misclassified cases, the results in Table 7 are almost identical to the main results in Table 5. The difference between Table 5 and Table 7 is not significant for any of the eight dependent variables. Therefore, misclassification of mothers is not impacting the results.

### **Heterogeneity by Child Gender and Maternal Preference Type**

As mentioned several times throughout this paper, there appears to be little overarching sex preference in Indonesia. Consistent with previous literature using data from Indonesia, there are no skewed sex ratios, mother stated sex preferences do not favor either sex, and there is no evidence that parents use sex-specific fertility stopping rules.<sup>17</sup> Still, checking for heterogeneity in the treatment effects by sex or maternal sex preference is important to confirm that the standardization of anthropometric measures is not providing misleading estimates. The first piece of supportive evidence comes from Table 5. Comparing the top and bottom panels, the inclusion of a male indicator does not affect the estimates. In addition, when full controls, including a male dummy, are included in the “With Controls” panel of Table 5, the estimated coefficient on the male dummy is never significantly different from zero. If the estimated effects for male and female children are different for standardized variables but not for the other health outcomes it might indicate that the main results are affected by the different relative heights and weights of Indonesian boys and girls. For example, consider the case where Indonesian boys have higher BMIs relative to

---

<sup>17</sup>Yamaguchi [1989], Clark [2000], Jayachandran [2014] discuss several implications of simple sex-specific fertility stopping rules. They should imply a negative relationship between the fraction of children that are male and family size and the sex-ratio at last birth parity should be male-biased. Neither of these relationships exist in the IFLS data. The correlation between whether a child is male and his/her family size is a precisely estimated zero. Further, the sex ratio at the last birth is not significantly different from .512 for families that have likely completed their fertility (those where the youngest child is at least five years of age) nor for families that may not have completed their fertility.

the male children in the reference distribution than Indonesian girls do relative to the female children in the reference distribution and the true treatment effect is zero. The main results could be positive simply because Indonesian boys are bigger than their female counterparts. The separate estimates by child gender would then indicate that the main estimates were driven entirely by male children. If, on the other hand, the sex-specific results are roughly equal for boys and girls, it is extremely unlikely that the main results were driven entirely by the standardization process. However, one caveat applies: there is no way to separately identify heterogeneity by sex and heterogeneity by maternal sex preference category. Therefore, these results should be interpreted as heterogeneity by both sex and maternal sex preference category.

Table 8 presents the average marginal effect of being born of the mother's preferred sex for both male children and female children. The results do not support that there is any important heterogeneity along this dimension. Only one of the eight columns estimates a treatment effect that is significantly different for boys (mothers who desired a boy) and girls (mothers who desired a girl). Also, this difference is for mother reported general child health, not for one of the anthropometric outcomes.

### **Saturated Control for Existing Household Composition**

Given how important the sex distribution among existing children is for predicting maternal gender preference, an important question is whether the results truly measure the effect of being born of your mothers more desired gender, or instead they simply reflect the direct effect of the existing sex distribution on child health outcomes. The data provide reason to doubt the this alternative explanation. First, consider the main estimates in Table 5. For each dependent variable results are displayed with and without controls. Among the included controls are a full set of indicators for the number of older male children in the household and the number of older female children in the household. If the sex distribution among existing children were driving the main results, I would expect the coefficient on

“Born of Mother’s Preferred Sex” to change with the inclusion of these variables as the estimate is purged of the source of the confounding variation. Instead, the estimates are nearly identical for each of the eight outcome variables, with and without controls.

However, the full set of controls contains much more than just the older sister/brother indicators. While unlikely, one might worry that the inclusion of just these indicators, and the omission of all other controls could still lead to starkly different treatment effect estimates. Further, while the older sister/brother indicators are fairly flexible, they do not completely saturate all possibilities for the existing household composition. That is, it may be the case that the effect of having one brother and one sister is not the same as the sum of the effect of having one brother and one sister. To allow for this possibility and to highlight the effect of just including controls for the existing household composition, Appendix B Table B1 estimates the effect of being born of the mother’s preferred sex on health outcomes with fully saturated controls for the existing household composition. This is done through the inclusion of a full set of number of older sister indicators, a full set of number of older brother indicators, and a full set of interactions between the sister and brother indicators. All other controls from the “With Controls” panel of Table 5 are omitted.

The results in Table B1 confirm that the main estimates are capturing variation in health outcomes generated by something other than the effect of the existing household composition. As was the case with Table 5, the inclusion of controls has no impact on the estimated effect of being born of the mother’s preferred sex. For all outcomes, the coefficients are statistically indistinguishable without any controls and with the fully saturated control for household composition.

What does this homogeneity imply for the interpretation of the main results? While Section V illustrates that the sex composition among existing children is the most powerful predictor of maternal sex preference, it says nothing about the direct effect of the sex composition on child health outcomes. That the inclusion of a fully saturated set of controls for household composition does not affect the treatment effects suggests that there is no



direct effect of the sex composition on child health outcomes. Instead, the only potential effect of the sex composition on health outcomes is an indirect effect: the sex composition potentially influences continued maternal fertility (and therefore family size), which then affects the health outcomes of the main child.

## **8 Mechanisms: Realized Gender and Subsequent Fertility**

As the predictions in the conceptual framework suggest, if continued maternal fertility is an important reason why children born of their mother's preferred sex receive are healthier, I would expect this result to be concentrated at parities where mothers still receive a sufficiently large marginal utility increase from an additional child. Specifically, the conceptual framework section suggests the effect should be more pronounced at lower birth parities.

To begin, Figure 2 explores at which birth parities the effect of subsequent fertility are concentrated. The figure plots a histogram of the ideal number of children that mothers in the sample desire. To avoid using reports that have been affected by the sex realizations of previous children, Figure 2 only uses information from women who had not yet begun their fertility. As suggested by Table 1, ideal fertility in Indonesia is fairly low. Roughly 80% of women desire three or fewer total children. Thus, for the vast majority of the mothers in the sample, Figure 2 suggests that they should begin facing a tradeoff between higher-than-desired total fertility and a potentially more balanced sex ratio around parity three.

Figure 3 investigates whether Proposition 5 is supported by the data. I do this by allowing the effect of being born of the mother's preferred sex to vary by birth parity. If the model is correct, Proposition 5 suggests that we should only see differences in health outcomes resulting from changes in fertility at parities one through three, or for children who have zero to two older siblings. The graphs in Figure 3 suggest this is very much the case

for BMI and weight-for-age. There is a drop in the estimated effect of being born of the mother's preferred sex between children who have two and children who have three older siblings. This corresponds with the parity at which most mothers would likely not continue fertility regardless of the sex composition among their existing children. Appendix Table B2 further investigates this theory by interacting the treatment indicator with an indicator for whether the child was born at parities one through three. The results indicate considerable heterogeneity in the treatment effect along this birth parity cutoff. For both BMI and Weight-for-age, the effect is large, positive, and statistically significant for children born at parities one through three. Conversely, the effect is negatively signed and never significant from zero at higher birth parities. Further, the p-value from a joint test of whether the coefficients are equal for children of parities one through three and those at higher parities, rejects the null hypothesis of no difference. The evidence supports the model prediction that the effects should be concentrated at low birth parities because at higher parities mothers will be unlikely to continue fertility regardless of the sex composition of existing children.

To continue exploring this potential mechanism, I directly estimate whether mothers reduce or delay later fertility if they have a child of their preferred sex. Then, making use of the fertility histories described in the data section, I calculate what fraction of the total effect on BMI and weight-for-age from Table 5 can be attributed to changes in subsequent maternal fertility.

Table 9 investigates whether mothers who have a child of their preferred sex are more or less likely to continue their fertility. I test for this by estimating (1) with two different dependent variables: an indicator for whether the mother continued her fertility and a count of the number of additional children she had. A significant, negative estimate of  $\beta$  indicates the mother is less likely to continue fertility if she had a child of her preferred sex. Consistent with this theory, the coefficient on "Born of Mother's preferred sex" is negative and significant in both columns. The point estimates indicate that mothers who had a child of their preferred sex were 5.8 percentage points less likely to have any additional children.

This effect is significant at the 5% level. Similarly, these mothers have fewer additional children. On average, mothers had .093 fewer additional children if they had a child of their preferred sex. As was the case in Table 5, the inclusion of the extensive set of controls has no impact on the estimated treatment effects. Both of these results are consistent with Proposition 3 from the model: mothers with a balanced sex ratio, those who had a child of their preferred sex, are less likely to continue fertility than those with a less balanced sex ratio.

There is a substantial literature, using both structural and reduced form methods, that finds birth spacing has important consequences for the health and human capital of the preceding and subsequent births as well as for the health of the mother [Palloni and Tienda, 1986, Bhalotra and Soest, 2008, DaVanzo et al., 2008, Buckles and Munnich, 2012]. Therefore, columns 3 and 4 of Table 9 display estimates from an exponential hazard model of the effect of having a child of the mother's preferred sex on the percentage change in the monthly hazard.<sup>18</sup> Column 3 includes all children, regardless of whether the mother continued her fertility, and top codes the interbirth interval to seventy-two months. Column 4 includes only children whose mothers had an additional child but similarly top codes the interval for those mothers who continued fertility to seventy-two months. Column 3 indicates that mothers who had a child of the sex they desired are 17% less likely to give birth to a subsequent child in each month. Column 4 suggests that much of the overall effect on birth spacing is driven by mothers who are not observed having another child; when limiting the sample to only mothers who have a subsequent child, the estimated effect decreases to .0395, implying roughly a 4% decrease in the likelihood of having an additional child each month. Still, even within the sample of mothers who continue their fertility, the interval following the birth of the main child is four months longer at the median for mothers who had a child of their preferred sex. Appendix Table C1 presents the non-parametric analogs—Kaplan-Meier cumulative failure plots—of Columns 3 and 4. Both panels clearly illustrate a

---

<sup>18</sup>The estimated effect of “Born of the Mother’s Preferred Sex” on the hazard rate is nearly identical using a Weibull or Gompertz distribution.

difference in the likelihood of having had another child beginning around 24 months after the birth of the main child. Whereas this difference persists indefinitely when including mothers who do not continue their fertility, the difference fades around 60 months when limiting the sample to mothers who continue their fertility.

Table 9 suggests that decreases in subsequent fertility could explain some of the main effects from Table 5. To assess how important changes in subsequent fertility are for explaining the changes in BMI and weight-for-age, I leverage the fact that mothers in Indonesia desire a balanced sex ratio among children. Section IV provides evidence that this is true for the mothers in my sample. The sex distribution among existing children explains a huge fraction of the variation in maternal desire for future sons and daughters. The higher the fraction of males among existing children, the lower the fraction of males the mother desires among future children. Further, the summary statistics in Table 1 indicate that mothers report wanting roughly the same number of male children as female children. Thus, stated maternal preferences for the women in the sample support the idea that mothers desire a balanced sex ratio.

If family size could be treated as exogenous to BMI and weight-for-age, a simple OLS regression could be used to identify the effect of family size on early health. Of course, the assumption that family size is exogenous to child health is quite strong. Children from larger families will differ from those in smaller families for a number of different reasons, many of which I will not be able to measure with IFLS data. As a result, OLS estimates of the effect of family size will likely be biased. To overcome this issue, I use the fertility history of all mothers in the IFLS to create two new variables: the total number of children born to each mother and an indicator for whether the sex of the mother's first child is different from the sex of her second child. Next, following Angrist and Evans [1998] and Cruces and Galiani [2007], I estimate the effect of family size on BMI and weight-for-age using an instrumental variables strategy where an indicator for whether the sexes of the mother's first two children are different (mixed sex indicator) is an excluded instrument for

the total number of children. I use a mixed sex indicator instead of two separate indicators for whether the first two children were male and whether the first two children were female as in Cruces and Galiani [2007]. I do this for two reasons. First, fertility is, on average, the same for mothers whose first two children were male and mothers whose first two children were female.<sup>19</sup> Second, my sample size is roughly one-twentieth of their sample. To ensure the estimates are precise enough to be informative I elect to pool across the first two male and first two female indicators.

Letting  $FS_m$  represent the total number of children for mother  $m$ ,  $SM_m$  be the mixed sex indicator, and suppressing all control variables from both equations, this two equation system can be written as:

$$y_i = \alpha + \theta FS_m + v_{im} \quad (2)$$

$$FS_m = \gamma + \phi SM_m + u_{im} \quad (3)$$

For the mixed sex indicator to be a valid instrument it must be the case that it is uncorrelated with  $v_{im}$ , the error term in the outcome equation, uncorrelated with  $u_{im}$ , the error term in the first stage, and that  $\phi \neq 0$ , so that whether the first two children are of different sexes has some effect on the total number of children born to the mother.

Ultimately, the first two assumptions are untestable. However, I argue that there are good reasons to believe that both hold in this context. First, as discussed throughout the paper, there are no skewed sex ratios in the sample. The ratio of males to females for the sample is approximately 105 to 100, a figure which is not significantly different from the biologically expected ratio of 104 males for every 100 females. This suggests there is no gender gap in mortality in Indonesia. Additionally, there is no evidence that parents practice sex-specific fertility stopping rules in Indonesia. Girls and boys have the same average

---

<sup>19</sup>To test this, I regress the total number of children for each mother on an indicator for whether her first two children were male, an indicator for whether her first two children were female, and an indicator for whether the first child was male. The test of whether the difference between the first two female and first two male coefficients is zero fails to reject the null (p-value .76).

family size in the sample and the sex ratio at the last birth (SRLB) is not significantly different from 104 males per 100 females. Together, this evidence supports the idea that the sex distribution is random in Indonesia and therefore  $SM_m$  is uncorrelated with  $u_{im}$ .

Whether  $\phi \neq 0$  is an empirical question. It amounts to the assumption that the first stage effect is not zero. Unfortunately, the main sample used in this paper is not large enough to precisely estimate a Two Stage Least Squares (2SLS) model with the same set of children. Therefore, I employ the Two Sample Two Stage Least Squares (TS2SLS) estimator proposed by Angrist and Krueger [1992]. This enables me to make use of the larger sample of households that have non-missing values for the instrument and the endogenous variable to more precisely estimate the first stage. To avoid weighting household with more children more than those with fewer children, I keep only one observation per mother for both the first stage and the second stage equations. For the second stage (reduced form) equation, I keep the oldest child for each mother in the data. The first stage also retains only one observation per mother from the first wave the mother participated in. Additionally, mothers must have at least two children and with non-missing values for gender. The oldest main sample child for each mother is then used to estimate the reduced form relationship between the instrument and the outcomes of interest. For the TS2SLS estimator to be appropriate, it must be true that one sample contain the instrument and the outcome measures, the other sample contain the instrument and the endogenous variable, and that the two samples be independent and drawn from the same population. In the case of the IFLS, not all children were selected for health measurement in the first three waves. From the roster of all children in the household, measured children were selected at random. If the randomly selected children was not also the child born immediately following the elicitation of the mother's sex preferences, then no mother-child observation from that household is used in the main analysis. However, I am able to use the fertility information, total number of children, and sex distribution among the oldest two children from that observation to help estimate the first stage equation. Crucially, what dictates whether the

mother-child observation can be used to estimate the reduced form, or just the first stage, is random sampling chance. Therefore, the two samples meet the requirements necessary for TS2SLS estimation.

Table 10 presents the results from the TS2SLS estimation. In both the reduced form and the first stage, survey wave dummies are included as controls. Including further controls for the sex of the child, the year of birth, month of birth, and birth order of the second stage child has no effect on the estimates. Column 1 displays the result of the first stage. As expected, mothers have significantly smaller families if their oldest two children are of the opposite sex. On average, they have .142 fewer children. This result is significant at the 1% level and has an F-statistic of 17.61.

Columns 2 and 3 present the TS2SLS estimates of the effect of family size on BMI and weight-for-age Z-scores, respectively. Standard errors are calculated based on 999 bootstrap repetitions.<sup>20</sup> An increase in family size has a negative and significant effect on the BMI Z-score and a negative, statistically insignificant effect on the weight-for-age Z-score of children. The point estimates suggest an extra child leads to a 1.056 standard deviation decrease in BMI and a .496 standard deviation decrease in weight-for-age in early childhood. These effects are quite large, between three and four times the size of the main results presented in Table 5.

With the TS2SLS estimates in hand, I now turn to calculating what fraction of the main results might plausibly be explained by changes in subsequent fertility. To estimate this fraction, I multiply the TS2SLS estimates of the effect of family size on each outcome by the average increase in fertility resulting from mothers not having a child of their preferred sex (Table 9). This product represents the average loss in BMI or weight-for-age that children likely experienced as a result of the increase in fertility. Next, I take the ratio

---

<sup>20</sup>More precisely, indexing the repetition number with  $r \in \{1, \dots, R\}$ , and letting  $\bar{\beta}_r$  be the mean estimate of the treatment effect across all repetitions, standard errors are calculated as  $se(\hat{\beta}) = \sqrt{\left(\frac{1}{R-1} \sum_{r=1}^R (\beta_r - \bar{\beta}_r)^2\right)}$ .

of this product to the corresponding effect from the main results in Table 5. The resulting fraction provides an estimate of what percent of the total treatment effect could be explained by changes in fertility. These fractions are shown in the bottom row of Table 10. I estimate that the fertility response by the mother explains nearly half of the treatment effect for BMI (42%) and almost thirty percent of the treatment effect for weight-for-age (28%).

## 9 Conclusion

I set out to investigate whether early childhood health is affected by whether children are born of their mother's preferred sex. Using an identification strategy that links a pre-conception measure of maternal preference to the realized sex and several early health indicators of each mother's next child, I find evidence that children born of the sex that their mother desired more are healthier along a number of observable dimensions. Children born of their mother's preferred sex have a significantly higher Body Mass Index, are heavier, and are less likely to suffer from wasting than children born of their mother's less preferred sex. In addition, these children experience fewer days of illness, fewer episodes of coughing and fever, and are reported as being generally healthier by their mothers. Together, these results strongly suggest that children who were born of their mother's preferred sex are receiving more resources from parents.

To help interpret the main results, Section VI explored what drives mothers to prefer that their next child be of a particular gender. In contrast to countries like India and China where son preference is widespread and likely driven by a variety of cultural, religious, and institutional factors, I find that nearly all of the variation in maternal desire for future sons is driven by the sex distribution among the mother's existing children. Mothers in Indonesia desire a balanced sex ratio, roughly equal numbers of males and females, among their children. This result is consistent with past literature on the relative lack of sex preference in Indonesia.



Leveraging the fact that mothers desire a mixed sex ratio, I estimate a TS2SLS model to test whether changes in fertility might be an important causal mechanism behind the main results. Using an indicator for whether a mother's first two children are of opposite sexes, I find that changes in fertility potentially explain 42% of the main effect on BMI and 28% of the main effect on weight-for-age, with the residual effect likely a result of reductions in parental consumption and the shifting of resources away from older children.

If parents have perfect information and they completely internalize the quantity-quality tradeoff, then this heterogeneous treatment is utility maximizing from their perspective. However, if parents are not fully cognizant of how continued fertility impacts the health and resources received by previous children, then providing parents with this information could reduce the disparity in continued fertility and the resulting gap in early childhood health. Documenting parents' beliefs about whether increased fertility negatively impacts the resources and health of existing children would provide critical guidance about whether disseminating information to parents is a viable policy tool.

Much of the policy discussion related to fertility and parental gender preference has focused on how decreases in fertility will affect sex-selection and sex-biased resource allocation. Taking as given the distribution of gender preferences, reductions in fertility are expected to lead to increases in sex-biased resource allocation and sex-biased mortality. The results of this paper suggest a policy instrument for affecting behavior in the opposite direction. If parental gender preferences are malleable, reducing the likelihood that parents prefer for their next child be of a particular sex will reduce total fertility. Focusing on interventions that increase parental indifference towards the sex of their next child has the potential to both lower total fertility and decrease disparities in early childhood health. Of course, the effectiveness of any such policy will depend critically on whether parental desire for the next child to be male or female can be influenced.

The results identified in this paper represent an important source of heterogeneity in the treatment of children. These differences in treatment create variation in real health

outcomes that are likely to persist into adulthood. Unrelated to health or skill endowments, children who are born of their mother's less preferred sex will face a gap in resources relative to children who were born of their mother's preferred sex. Identifying the optimal policies to close this gap in early conditions remains a question for future research.

## 10 Tables

**TABLE 1: Summary Statistics**

	mean	sd	min	p50	max	count
<i>Panel A</i>						
Subsequent Child born of Mother's Preferred Sex	0.468	0.499	0.000	1.000	1.000	1517
Mother Prefers Girls	0.219	0.413	0.000	0.000	1.000	3490
Mother Prefers Boys	0.216	0.412	0.000	0.000	1.000	3490
Mother's Number of Desired Additional Children	1.209	1.091	0.000	1.000	12.000	3490
Mother's Total Ideal Number of Daughters	1.359	0.956	0.000	1.000	7.000	3753
Mother's Total Ideal Number of Sons	1.413	0.991	0.000	1.000	8.000	3753
Male	0.507	0.500	0.000	1.000	1.000	3755
Age in Months	41.241	26.845	0.000	35.000	98.000	3586
Number of Older Brothers	0.932	1.049	0.000	1.000	7.000	3695
Number of Older Sisters	0.876	1.021	0.000	1.000	7.000	3694
Years of Schooling Completed by Mother	8.228	4.364	0.000	7.000	20.000	3379
Mother is Muslim	0.903	0.296	0.000	1.000	1.000	3380
Mother is Christian	0.048	0.214	0.000	0.000	1.000	3380
Mother is Hindi	0.046	0.210	0.000	0.000	1.000	3380
Mother Ever Worked During Past Year	0.469	0.499	0.000	0.000	1.000	3381
Mother's Primary Activity: Working	0.391	0.488	0.000	0.000	1.000	2498
Mother's Primary Activity: Housekeeping	0.596	0.491	0.000	1.000	1.000	2498
<i>Outcomes</i>						
Illness (4 Weeks): Days Primary Activity Missed	1.666	3.043	0.000	0.000	28.000	3743
Illness (4 Weeks): Days Bedridden	0.420	1.604	0.000	0.000	28.000	3743
BMI Z-Score	-0.349	1.661	-5.000	-0.386	5.000	3540
Weight-for-age Z-Score	-1.181	1.434	-6.000	-1.257	5.000	3578
Height-for-age Z-Score	-1.540	1.672	-6.000	-1.620	6.000	3545
Child is Wasted (WHO Definition)	0.111	0.314	0.000	0.000	1.000	3540
Child is Thin (WHO Definition)	0.264	0.441	0.000	0.000	1.000	3578
Child is Stunted (WHO Definition)	0.389	0.488	0.000	0.000	1.000	3545
<i>Panel B</i>						
Child is Male	0.514	0.500	0.000	1.000	1.000	24348
Birth Order	2.547	1.643	1.000	2.000	17.000	24348
Number of Children	4.112	2.071	2.000	4.000	17.000	24348
Oldest Two Sibs Sex Mixed	0.468	0.499	0.000	0.000	1.000	24348
Age in Years	8.034	7.366	0.000	6.000	53.000	24348

Notes: Data taken from the four waves of the Indonesian Family Life Survey. Sample descriptions contained in the main text.

**TABLE 2: Summary Statistics by Maternal Preference Category**

	Mean NGP (1)	Mean SP (2)	Mean DP (3)	Diff 1 and 2 (4)	P-Val (5)	Diff 1 and 3 (6)	P-Val (7)	Diff 2 and 3 (8)	P-Val (9)
<i>Household and Child Characteristics</i>									
Mother's Number of Desired Additional Children	1.035	1.472	1.397	-0.436	<b>0.000</b>	-0.362	<b>0.000</b>	0.074	0.110
Mother's Ideal Number of Daughters	1.354	1.419	1.507	-0.065	0.127	-0.153	<b>0.000</b>	-0.088	<b>0.029</b>
Mother's Ideal Number of Sons	1.415	1.559	1.471	-0.144	<b>0.001</b>	-0.056	0.183	0.088	<b>0.038</b>
Number of Older Brothers	1.043	0.260	1.343	0.783	<b>0.000</b>	-0.300	<b>0.000</b>	-1.083	<b>0.000</b>
Number of Older Sisters	0.981	1.252	0.248	-0.271	<b>0.000</b>	0.733	<b>0.000</b>	1.004	<b>0.000</b>
Male	0.510	0.470	0.533	0.040	<b>0.062</b>	-0.023	0.275	-0.063	<b>0.014</b>
Age in Months	40.987	41.509	39.896	-0.522	0.656	1.090	0.349	1.613	0.241
Years Schooling Completed by Mother	8.087	8.483	8.630	-0.396	<b>0.046</b>	-0.543	<b>0.006</b>	-0.147	0.508
Mother is Muslim	0.910	0.877	0.908	0.033	<b>0.014</b>	0.002	0.877	-0.031	<b>0.060</b>
Mother is Christian	0.046	0.056	0.044	-0.010	0.304	0.003	0.784	0.013	0.283
Mother is Hindi	0.040	0.064	0.048	-0.024	<b>0.011</b>	-0.008	0.342	0.016	0.202
Mother Reported Health of Child Z-Score	-0.015	-0.050	0.074	0.036	0.406	-0.089	<b>0.037</b>	-0.124	<b>0.014</b>
Surveyor Reported Health of Child Z-Score	-0.025	0.032	0.033	-0.056	0.203	-0.057	0.195	-0.001	0.983
Illness (4 Weeks): Days Primary Activity Missed	1.612	1.668	1.823	-0.056	0.658	-0.211	0.107	-0.155	0.356
Illness (4 Weeks): Days Bedridden	0.398	0.363	0.511	0.034	0.585	-0.113	0.104	-0.148	0.103
Hemoglobin Level (g/dl)	11.292	11.344	11.322	-0.052	0.437	-0.030	0.664	0.023	0.779
Anemic	0.404	0.394	0.373	0.010	0.701	0.031	0.212	0.021	0.469
BMI Z-Score	-0.341	-0.368	-0.365	0.026	0.720	0.023	0.748	-0.003	0.975
Weight-for-age Z-Score	-1.174	-1.191	-1.169	0.017	0.788	-0.005	0.939	-0.022	0.770
Height-for-age Z-Score	-1.539	-1.517	-1.515	-0.023	0.756	-0.024	0.745	-0.001	0.988
Child is Wasted (WHO Definition)	0.118	0.096	0.110	0.022	0.113	0.009	0.530	-0.013	0.413
Child is Thin (WHO Definition)	0.266	0.248	0.269	0.019	0.332	-0.002	0.899	-0.021	0.359
Child is Stunted (WHO Definition)	0.387	0.384	0.393	0.003	0.876	-0.006	0.786	-0.009	0.722
Mother's BMI-for-age Z-Score	0.017	0.009	-0.002	0.008	0.858	0.018	0.671	0.010	0.832

Notes: See Table 1. The NGP column (1), contains information on children born to mothers who desired the same number of future male and female children before their birth. The SP column (2), displays means for children born to mothers who desired more future sons than daughters before their birth. The DP column (3), displays means for those children born to mothers who desired more future daughters than sons before their birth. Columns 4 and 5 show the differences-in-means and the P-value from a test of whether there is no difference. Columns 6 and 7 do the same for the comparison between NGP and SP children, and columns 8 and 9 the same for SP and DP children. The child health Z-scores are constructed so that lower values imply healthier children.

**TABLE 3: Summary Statistics by Preferred or Less Preferred Sex**

	Mean LP (1)	Mean MP (2)	Diff (3)	P- Value (4)	N Obs LP (5)	N Obs MP (6)
<i><u>Household and Child Characteristics</u></i>						
Mother Prefers Girls	0.504	0.500	0.004	0.866	807	710
Mother Prefers Boys	0.496	0.500	-0.004	0.866	807	710
Mother's Number of Desired Additional Children	1.442	1.424	0.018	0.693	807	710
Mother's Ideal Number of Daughters	1.481	1.442	0.039	0.340	807	710
Mother's Ideal Number of Sons	1.518	1.511	0.007	0.875	807	710
Male	0.504	0.500	0.004	0.866	807	710
Age in Months	40.922	40.396	0.525	0.703	767	681
Household Size	5.954	5.968	-0.013	0.910	807	710
Household Lives in an Urban Area	0.475	0.465	0.010	0.703	807	710
Number of Older Brothers	0.801	0.807	-0.006	0.901	799	704
Number of Older Sisters	0.763	0.730	0.034	0.455	799	703
Mother's Age	26.563	26.882	-0.319	0.225	807	709
Years of Schooling Completed by Mother	8.515	8.600	-0.085	0.703	730	648
Mother Ever Worked During Past Year	0.474	0.488	-0.014	0.613	732	650
Mother is Muslim	0.885	0.902	-0.016	0.324	731	650
Mother is Christian	0.059	0.040	0.019	0.109	731	650
Mother is Hindi	0.056	0.055	0.001	0.955	731	650
<i><u>Child Outcomes</u></i>						
Mother Reported Health of Child Z-Score	0.049	-0.030	0.079	0.119	805	708
Surveyor Reported Health of Child Z-Score	0.024	0.042	-0.018	0.723	766	678
Illness (4 Weeks): Days Primary Activity Missed	1.954	1.510	0.444	<b>0.008</b>	804	708
Illness (4 Weeks): Days Bedridden	0.497	0.370	0.127	0.163	805	708
Hemoglobin Level (g/dl)	11.369	11.292	0.077	0.341	579	506
Anemic	0.380	0.387	-0.007	0.803	579	506
BMI Z-Score	-0.470	-0.248	-0.222	<b>0.010</b>	761	667
Weight-for-age Z-Score	-1.253	-1.097	-0.156	<b>0.035</b>	767	678
Height-for-age Z-Score	-1.505	-1.528	0.022	0.799	761	669
Child is Wasted (WHO Definition)	0.117	0.087	0.030	<b>0.063</b>	761	667
Child is Thin (WHO Definition)	0.272	0.242	0.031	0.185	767	678
Child is Stunted (WHO Definition)	0.377	0.401	-0.023	0.364	761	669

Notes: See Table 1. The LP column (1) presents means for children born of their mothers ex-ante less preferred sex. The MP column (2) presents means for children born of their mothers more preferred sex. Column 3 displays the difference-in-means between columns 1 and 2, column 4 displays the associated P-value from a test of whether there is no difference between the two columns, column 5 shows the number of LP children with non-missing values, and column 6 shows the number of MP children with non-missing values. The child health Z-scores are constructed so that lower values imply healthier children.

**TABLE 4: Explaining The Desired Male Ratio**

<i>Panel A:</i>	Full Controls	No Controls
<i>Variable</i>	(1)	(2)
Current Male Ratio Among Children (Boys/Total)	-0.735*** (0.016)	-0.735*** (0.016)
Average Bridewealth Province of Residence (Thousands of Rupiah)	0.000 (0.000)	
Fraction of Women in Province Moved from Village Upon Marriage	-0.082 (0.137)	
Muslim	-0.224 (0.198)	
Hindu	-0.175 (0.201)	
Christian	-0.193 (0.201)	
Mother Attended Primary School	0.072** (0.035)	
Mother Attended Junior High School	0.085** (0.038)	
Mother Attended Senior High School	0.093** (0.038)	
Mother's Earnings Last Year (Hundreds of Thousands of Rupiah)	-0.001*** (0.000)	
Constant	1.304*** (0.342)	0.877*** (0.011)
Observations	1,419	1,419
R-squared	0.603	0.586
Chi-2 Stat: Coeff on Male Ratio is Equal in (1) and (2)		0.0110
P-value		0.917
<b><i>Panel B: Predicting Maternal Son Preference</i></b>		
	% Correctly	
	Predicted	se
Full Controls	84.1	0.007
No Controls	83.4	0.008
Observations		2397
Difference in % Correctly Predicted		0.007
P-Value: Difference is Zero		0.457

Notes: See Table 1. Panel A presents the results from OLS regressions of the mother desired fraction male among future children (desired future male children/total desired future children) on the fraction of existing children that are male. Column 1 also includes the full set of controls discussed in the text. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1. Panel B displays the accuracy of a linear probability model (LPM) for predicting maternal son preference using only the fraction of existing children that are male or the fraction of existing children that are male as well as the full set of controls included in Panel A. Mothers are designated as having a "predicted son preference" if the LPM predicts a probability of son preference of over .5 and "not predicted to have a son preference" otherwise.

**TABLE 5: Main Results With and Without Controls**

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	BMI Z-Score	Weight-for-age Z-Score	Mother Reported Health of Child Z-Score (Lower is Healthier)	Days of Primary Activity Missed by Child Due to Illness	Days Spent in Bed by Child Due to Illness	Child is Wasted (WHO Definition)	Had Cough	Had Fever
<b><u>Without Controls</u></b>								
Born of Mother's Preferred Sex	0.236*** (0.086)	0.164** (0.075)	-0.091* (0.052)	-0.499*** (0.170)	-0.161* (0.092)	-0.032** (0.016)	-0.073*** (0.026)	-0.056** (0.026)
Constant	-0.474*** (0.059)	-1.252*** (0.047)	0.061* (0.036)	1.978*** (0.135)	0.512*** (0.074)	0.118*** (0.012)	0.462*** (0.018)	0.450*** (0.018)
Observations	1,414	1,430	1,433	1,432	1,433	1,414	1,433	1,433
<b><u>With Controls</u></b>								
Born of Mother's Preferred Sex	0.239*** (0.089)	0.151** (0.074)	-0.086 (0.053)	-0.437** (0.171)	-0.153* (0.092)	-0.029* (0.017)	-0.070*** (0.026)	-0.048* (0.026)
Constant	-132.884 (116.689)	-87.472 (93.637)	60.276 (51.420)	-873.427** (422.030)	-88.465 (98.336)	8.582 (18.940)	-25.087 (39.336)	42.598 (38.581)
Observations	1,414	1,430	1,433	1,432	1,433	1,414	1,433	1,433
P-Value: Coefficients Are the Same	0.913	0.679	0.729	0.274	0.776	0.623	0.748	0.325

Notes: The sample includes children conceived immediately following the elicitation of their mothers preferences over subsequent child gender. In addition, their mothers must have either preferred male children or female children, they could not have been indifferent to either sex. Both panels present the results of OLS regressions. The without controls panel includes only the "Born of Mother's Preferred Sex" indicator and a constant. The with controls panel includes additional controls for whether the mother preferred to have a female child, whether the child is male, the child's age-in-months, the mother's age, a full set of indicators for the number of older brothers and the number of older sisters each child has, total household size, year of birth dummies, month of birth dummies, province dummies, linear time trends that are allowed to differ by province, an indicator for whether the household lives in an urban area, and dummies for the wave of the survey the observation comes from. In addition, each column shows the P-value from a test of whether the coefficient on "Born of Mother's Preferred Sex" is equal in the two panels. Robust standard errors clustered at the mother level are in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.

**TABLE 6**  
**Selective Abortion, Selective Mortality and Early Conditions**

	Mother had a Non-Live Birth Between Waves			Any live births that died between waves	Birth Weight (Kg)	Gestational Length (Weeks)
	(1)	(2)	(3)	(4)		
Born of Mother's Preferred Sex	0.009 (0.010)	0.012 (0.011)	-0.017 (0.035)	0.029 (0.104)		
Constant	0.036*** (0.007)	0.043*** (0.007)	3.278*** (0.024)	39.116*** (0.070)		
Observations	1,433	1,433	971	1,416		

Notes: The sample includes children conceived immediately following the elicitation of their mothers preferences over subsequent child gender. In addition, their mothers must have either preferred male children or female children, they could not have been indifferent to either sex. All columns present the results of OLS regressions with no controls. Columns 1 and 2 consider two outcomes intended to measure the likelihood that mothers experienced any pre-natal or early childhood mortality between survey waves. The dependent variable in column 1 is an indicator variable equal to one of the mother had a miscarriage between stating her sex preferences and the birth of the child included in the main analysis. The dependent variable in column 2 is an indicator variable for whether the mother gave birth to a living child that subsequently died between stating her sex preferences and the birth of the child included in the main analysis. Columns 3 and 4 consider two dependent variables intended to capture the initial health of the child used in the main analysis. Column 3 tests whether children that were born of their mother's preferred sex are heavier at birth. Column 4 tests whether children that were born of their mother's preferred sex had longer gestations. Robust standard errors clustered at the mother level are in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.



**TABLE 7**  
**Limiting the Sample to Mothers That Desire One Gender Only**

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	BMI Z-Score	Weight-for-age Z-Score	Mother Health of Child Z-Score (Lower is Healthier)	Days of Primary Activity Missed by Child Due to Illness	Days Spent in Bed by Child Due to Illness	Child is Wasted (WHO Definition)	Had Cough	Had Fever
<b><u>Full Sample Without Controls</u></b>								
Born of Mother's Preferred Sex	0.236*** (0.086)	0.164** (0.075)	-0.091* (0.052)	-0.499*** (0.170)	-0.161* (0.092)	-0.032** (0.016)	-0.073*** (0.026)	-0.056** (0.026)
Constant	-0.474*** (0.059)	-1.252*** (0.047)	0.061* (0.036)	1.978*** (0.135)	0.512*** (0.074)	0.118*** (0.012)	0.462*** (0.018)	0.450*** (0.018)
Observations	1,414	1,430	1,433	1,432	1,433	1,414	1,433	1,433
<b><u>Single Sex Desired Only Without Controls</u></b>								
Born of Mother's Preferred Sex	0.239** (0.095)	0.174** (0.081)	-0.098* (0.056)	-0.388** (0.184)	-0.180* (0.103)	-0.033* (0.017)	-0.076*** (0.028)	-0.070** (0.028)
Constant	-0.459*** (0.064)	-1.246*** (0.051)	0.052 (0.038)	1.934*** (0.143)	0.530*** (0.084)	0.117*** (0.013)	0.458*** (0.019)	0.457*** (0.019)
Observations	1,223	1,238	1,241	1,240	1,241	1,223	1,241	1,241

Notes: The sample includes children conceived immediately following the elicitation of their mothers preferences over subsequent child gender. In addition, to be included in the results displayed in the bottom panel, their mothers must have either desired only male children or only female children, they could not have wanted some additional children of each sex. Robust standard errors clustered at the mother level are in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.

**TABLE 8**  
**Heterogeneity by Gender**

	Days of							
	Mother Reported Health of Child Z-Score		Primary Activity Missed by Child Due to Illness		Days Spent in Bed by Child Due to Illness		Child is Thin (WHO Definition)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<b><u>Girls</u></b>								
Born of Mother's Preferred Sex	0.242** (0.119)	0.199* (0.109)	0.023 (0.072)	-0.314 (0.252)	-0.026 (0.128)	-0.021 (0.022)	-0.076** (0.037)	-0.053 (0.037)
<b><u>Boys</u></b>								
Born of Mother's Preferred Sex	0.230* (0.126)	0.129 (0.103)	-0.204*** (0.074)	-0.681*** (0.229)	-0.295** (0.131)	-0.042* (0.023)	-0.069* (0.037)	-0.060 (0.037)
Observations	1,414	1,430	1,433	1,432	1,433	1,414	1,433	1,433
P-Value Effects Equal For Boys and Girls	0.946	0.641	0.0270	0.282	0.140	0.507	0.889	0.892

Notes: The sample includes children conceived immediately following the elicitation of their mothers preferences over subsequent child gender. In addition, their mothers must have either preferred male children or female children, they could not have been indifferent to either sex. Both panels present the results of OLS regressions. The girls panel for female children and the boys panel for male children. In addition, each column shows the P-value from a test of whether the coefficient on "Born of Mother's Preferred Sex" is equal in the two panels. Robust standard errors clustered at the mother level are in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.

**TABLE 9:  
Maternal Gender Preference, Realized Gender and Subsequent Fertility**

	Any Younger Siblings (1)	Number of Younger Siblings (2)	Months Until Next Birth (3)	Months Until Next Birth (Intensive Margin) (4)
<u><i>Without Controls</i></u>				
Born of Mother's Preferred Sex	-0.058** (0.026)	-0.093** (0.039)	-0.170* (0.079)	-0.0395 (0.035)
<u><i>With Controls</i></u>				
Born of Mother's Preferred Sex	-0.064*** (0.023)	-0.103*** (0.034)	-0.222** (0.077)	-0.0337 (0.037)
Observations (Children)	1,433	1,433	1,433	1,433
P-Value: Coefficients Are the Same	0.673	0.623	0.308	0.816
Median Difference in Birth Spacing (Months)				4

Notes: The sample includes children conceived immediately following the elicitation of their mothers preferences over subsequent child gender. In addition, their mothers must have either preferred male children or female children, they could not have been indifferent to either sex. Columns 1 and 2 of both panels present the results of OLS regressions. The without controls panel includes only the "Born of Mother's Preferred Sex" indicator and a constant. The with controls panel includes additional controls for whether the mother preferred to have a female child, whether the child is male, the child's age-in-months, the mother's age, a full set of indicators for the number of older brothers and the number of older sisters each child has, total household size, year of birth dummies, province dummies, linear time trends that are allowed to differ by province, an indicator for whether the household lives in an urban area, and dummies for the wave of the survey the observation comes from. In addition, each column shows the P-value from a test of whether the coefficient on "Born of Mother's Preferred Sex" is equal in the two panels. Columns 3 and 4 present the results of a maximum likelihood exponential survival analysis. In both columns, the estimates represent the percentage change in the risk of having an additional child, per month, if the previous child was born of the mother's preferred sex as opposed to being born of the mother's less preferred sex. Column 3 topcodes birth spacing to be 72 months and includes mothers who do not have an additional child. Column 4 also topcodes birth spacing to be 72 months but the sample includes only mothers who have an additional child. The difference between the median birth spacing for mothers who have an additional child is displayed at the bottom of column 4. Estimates using Weibull and Gompers distributions are nearly identical. Robust standard errors clustered at the mother level are in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.


**TABLE 10: Family Size, Maternal Gender Preference and Continued Fertility**

VARIABLES	First Stage:		
	Estimate on Family Size (1)	BMI Z-Score (2)	Weight-for-age Z-Score (3)
Oldest Two Children of Mother Are of The Opposite Sex	-0.142*** (0.034)		
Family Size		-1.056* (0.566)	-0.496 (0.426)
Constant	4.069 (0.029)		
Observations	7,812	2,744	2,744
R-squared	0.169		
First Stage F-Statistic on Excluded Instrument	17.61		
Fraction of Total Effect Explained by Continued Fertility		0.42	0.28

Notes: Table 10 presents coefficients from Two Sample Two Stage Least Squares (TS2SLS) estimates of the effect of family size on BMI-for-age Z-score and Weight-for-age Z-score. The excluded instrument in the first stage is an indicator for whether the oldest two children born to the mother were of different sexes. Column 1 presents the first stage estimate on the excluded instrument, the first stage r-squared, and the F-statistics from a test of whether the excluded instrument is zero in the first stage. Both stages include controls for the survey wave the observation is from. At the bottom of columns 2 and 3, the table shows the fraction of the main results (from Table 5) that can be explained by continued fertility by the mother. This is calculated by multiplying the estimated effect of Family Size on each outcome from Table 10 by the estimated increase in fertility from Table 9 (column 2). The percent explained is then given by the ratio of this product to the corresponding effect from the main results in Table 5. Standard errors in parentheses. Standard errors for the TS2SLS estimates are bootstrapped based on 999 repetitions. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.


# 11 Figures

Figure 1  
Likelihood Child is Male By Birth Parity



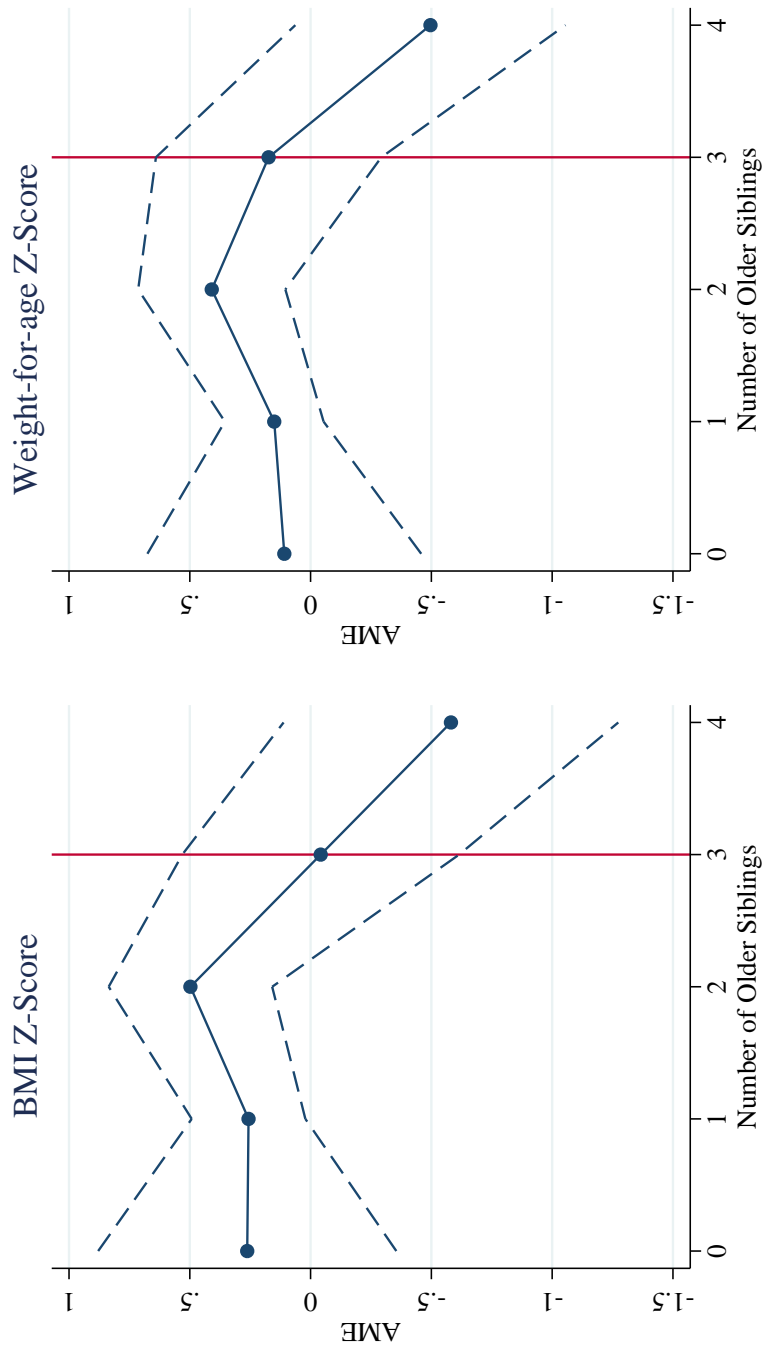
The fraction of children who are male at all parities less than 8 is .514. The test of whether this fraction is significantly different from the biologically expected .512 yields a P-value of less than .001.

Figure 2: Mother's Desired Number of Children



Histogram of the total number of children desired by mothers in the sample who have no prior children. The sample is limited to mothers with no children to avoid the possibility that past child gender realizations have influenced the total number of children desired. The vertical line is between three and four desired children. 78% of these mothers desire three or fewer children.

Figure 3: Average Marginal Effects  
By Birth Parity



Average marginal effects of treatment at different birth parities. The number of older siblings is top coded to four. 95% Confidence Intervals represented by dashed lines.

## 12 Appendix A: Model

To help illustrate the decisions mothers face, this section outlines a simple utility maximization problem. In keeping with the empirical findings in this and previous papers set in contemporary Indonesia, the model assumes that there are no systematic individual, cultural, or economic biases that lead mothers to favor male children over female children. Specifically, the model assumes that mothers face the same price of investment for male and female children, that mothers allocate resources equally across all existing children, and that mothers value the quality and count of male and female children equally. Despite this lack of overarching gender bias, the model is able to produce differences in child quality by gender. Instead of consistently favoring children of one sex, these differences reflect heterogeneity in received resources and health according to whether the child's sex matched the sex the mother desired more before their birth.

Mothers receive utility from the number of children they have  $n$ , the average health or quality of existing children  $\bar{q}$ , and the sex composition of existing children  $\frac{m}{n}$  where  $m$  is the number of existing male children. The existing sex composition enters the mother's utility function as a distance from her ideal ratio,  $\gamma$ . As mentioned above, much of the subsequent discussion will assume that  $\gamma = \frac{1}{2}$ , so that mother's desire a balanced sex ratio among children. Further, by maintaining that only average child quality enters the mother's utility function and that quality for each child  $q$  is simply defined as the investment  $i$  in that child in a period, I assume away the possibility that mothers allocate resources unequally across existing children: mothers will set investment (quality) to be equal across all children. I do this to show that an unequal allocation of resources among existing children is not needed to produce differences in health between children that were born of their mother's preferred sex and those that were born of their mother's less preferred sex.<sup>21</sup> Because the only uncertainty in the model comes from the lack of control over the gender of future

---

<sup>21</sup>To the extent that this differential allocation across existing children is important, it should reinforce the effects highlighted by this model.



children, I focus on the static problem faced by mothers who had a child during the previous period. After observing the sex of their most recent child, these mothers elect whether to have an additional child or not  $c \in \{0, 1\}$ . Mothers believe that their new child will be male with probability  $\frac{1}{2}$ , they face a price of investment  $p_i$ , and they receive exogenous resources  $y$  to split evenly across existing children.

I assume the mother's utility function is additively separable in  $n$ ,  $\bar{q}$ , and  $d(\cdot)$ , so the per period utility can be written as:

$$u\left(n, \bar{q}, d\left(\frac{m}{n} - \gamma\right)\right) = f(n) + \alpha g(\bar{q}) - \theta d\left(\frac{m}{n} - \gamma\right) \quad (4)$$

With no opportunity to smooth consumption across periods, mothers face the per period budget constraint:

$$y = p_i n \quad (5)$$

I assume that the net benefit from an additional child, absent any consideration of the sex ratio, is increasing and concave in  $n$ . That is,  $b(n) \equiv f(n) + \alpha g(\bar{q})$  is such that:

$$b'(n) > 0, b'' < 0 \quad (6)$$

Similarly, I assume the distance function is increasing and concave in the distance between the existing sex ratio and the ideal sex ratio ( $\gamma$ ):

$$d' > 0, d'' < 0 \quad (7)$$

Assuming that  $d(\cdot)$  is concave amounts to assuming that mothers care more about changes in the sex ratio close to their ideal ratio than they do about changes in the sex ratio far away from their ideal ratio. For example, it implies that, ignoring utility differences resulting from changes in family size, a mother who desires an equal number of boys and girls value a move from four boys and two girls to three boys and three girls more than she would a move from five boys and one girl to four boys and two girls.

The scale factors  $\alpha$  and  $\theta$  capture the relative importance of average child quality and

the distance from the mother's ideal sex ratio as compared to the number of children the mother has.

The discrete choice problem faced by the mother can be written as:

$$\begin{aligned} & \max_c \{u_{c=0}^*, u_{c=1}^*\} \\ & = \max_c \left\{ f(n) + \alpha g\left(\frac{y}{p_i n}\right) - \theta d\left(\frac{m}{n} - \gamma\right), f(n+1) + \alpha g\left(\frac{y}{p_i(n+1)}\right) - \frac{1}{2}\theta d\left(\frac{m}{n+1} - \gamma\right) - \frac{1}{2}\theta d\left(\frac{m+1}{n+1} - \gamma\right) \right\} \end{aligned}$$

Where the second line uses that  $\bar{q}_n = \bar{l}_n = i = \frac{y}{p_i n}$ .

This formulation of the mother's problem is still powerful enough to generate variation in outcomes between children that are born of their mother's preferred sex and those that are born of their mothers less preferred sex. In particular, the model suggests that at lower birth parities mothers will be less likely to continue their fertility if the sex ratio among existing children is equal to their ideal sex ratio  $\gamma$ . However, at higher birth parities mothers are unlikely to continue fertility regardless of the sex ratio among their existing children.

### Predictions for Stated Maternal Sex Preferences

To begin, consider how the model predicts that the mother's stated sex preferences will behave. Mothers who desire an additional child will prefer to have that child be of the sex that maximizes their utility in the next period. For a mother with  $m = \tilde{m}$ ,  $n = \tilde{n}$ , and  $\gamma = \tilde{\gamma}$ , this amounts to selecting the max from the following:

$$\left\{ u_{m+1}\left(\tilde{n}+1, \bar{q}_{\tilde{n}+1}, d\left(\frac{\tilde{m}+1}{\tilde{n}+1} - \tilde{\gamma}\right)\right), u_{f+1}\left(\tilde{n}+1, \bar{q}_{\tilde{n}+1}, d\left(\frac{\tilde{m}}{\tilde{n}+1} - \tilde{\gamma}\right)\right) \right\}$$

Using the functional form assumption outlined above, this simplifies to selecting the max from:

$$\left\{ -d\left(\frac{\tilde{m}+1}{\tilde{n}+1} - \tilde{\gamma}\right), -d\left(\frac{\tilde{m}}{\tilde{n}+1} - \tilde{\gamma}\right) \right\}$$

**Proposition 1** Mothers will prefer that their next child be of the gender that moves them closer to a balanced sex ratio.

Unsurprisingly, the mother will state a preference for a child of whichever sex brings her

closer to the ideal sex ratio  $\tilde{\gamma}$ . If mothers desire a balanced sex ratio (i.e.  $\tilde{\gamma} = \frac{1}{2}$ ), mothers will state a preference for whichever gender child they have fewer of. In this way, even without different prices of investment for male and female children, different marginal utilities for female child quality and male child quality (i.e.  $u'_{\bar{q}_{male}} \neq u'_{\bar{q}_{female}}$  with  $\frac{m}{n} = \frac{1}{2}$ ), or mothers who desire unbalanced sex ratios ( $\gamma \neq \frac{1}{2}$ ), stated preferences will differ across mothers as a result of the realized gender of previous children.

### Predictions for Continued Fertility and Resource Allocation

Mothers who prefer to have a male child next will not necessarily treat their next child better if it is a male.<sup>22</sup> However, even when assuming that mothers can't favor some of their existing children by dividing up resources unequally, allowing future fertility to depend on the sex of the most recent child produces the prediction that, at lower birth parities, children that are born of the sex their mother desired before their conception will receive more resources and have higher average quality. The following are the implications of the model assuming  $\gamma = \frac{1}{2}$  so that mothers desire a balanced sex ratio among children.

**Proposition 2** If a mother decides not to have a child in a period, she will also elect not to have a child in subsequent periods.

Consider a mother with  $m$  male children and  $n$  total children. If the mother elects not to have a child in period  $t$  it implies:

$$u_{t,c=1}^* \leq u_{t,c=0}^*$$

$$\Rightarrow \frac{1}{2} \left\{ f^{(n+1)} + \alpha_g \left( \frac{y}{p_i(n+1)} \right) - \theta_d \left( \frac{m+1}{n+1} - \frac{1}{2} \right) \right\} + \frac{1}{2} \left\{ f^{(n+1)} + \alpha_g \left( \frac{y}{p_i(n+1)} \right) - \theta_d \left( \frac{m}{n+1} - \frac{1}{2} \right) \right\} \leq f^{(n)} + \alpha_g \left( \frac{y}{p_i n} \right) - \theta_d \left( \frac{m}{n} - \frac{1}{2} \right)$$

As the only randomness in the model is a result of the unknown gender of the next child, the decision problem for a mother who elects not to have a child is exactly the same

---

<sup>22</sup>To see this, consider changing  $\gamma$  and allowing  $\bar{q}_{male}$  and  $\bar{q}_{female}$  to enter the mother's utility function separately and asymmetrically. With  $\gamma < \frac{1}{2}$ , and  $MU_{\bar{q}_{male}} > MU_{\bar{q}_{female}}$ , mothers will desire more female children but they will invest more in male children.

in the next period. Therefore, the mother will prefer  $u_{s,c=0}^*$  to  $u_{s,c=1}^*$  for all  $s \geq t$ .

**Proposition 3** Mothers are at least as likely to continue fertility when they are farther away from their target ratio.

Consider the decision faced by a mother who is under her target ratio by one  $\frac{m}{n} < \frac{1}{2}$  versus she is exactly at her target ratio  $\frac{m+1}{n} = \frac{1}{2}$ .

In the former case ( $\frac{m}{n} < \frac{1}{2}$ ) the mother compares her expected utility from having an additional child:

$$u_{c=1}^1 = \frac{1}{2} \left\{ f(n+1) + \alpha g \left( \frac{y}{p_i(n+1)} \right) - \theta d \left( \frac{m+1}{n+1} - \frac{1}{2} \right) \right\} + \frac{1}{2} \left\{ f(n+1) + \alpha g \left( \frac{y}{p_i(n+1)} \right) - \theta d \left( \frac{m}{n+1} - \frac{1}{2} \right) \right\}$$

$$u_{c=1}^1 = f(n+1) + \alpha g \left( \frac{y}{p_i(n+1)} \right) - \frac{1}{2} \theta d \left( \frac{m+1}{n+1} - \frac{1}{2} \right) - \frac{1}{2} \theta d \left( \frac{m}{n+1} - \frac{1}{2} \right)$$

To her utility from not having an additional child:

$$u_{c=0}^1 = f(n) + \alpha g \left( \frac{y}{p_i n} \right) - \theta d \left( \frac{m}{n} - \frac{1}{2} \right)$$

In other words, letting  $b_{diff}(n+1, n) \equiv b(n+1) - b(n)$ , the mother will elect to continue fertility iff:

$$u_{c=1}^1 \geq u_{c=0}^1$$

$$\implies b_{diff}(n+1, n) - \frac{1}{2} \theta \left( d \left( \frac{m+1}{n+1} - \frac{1}{2} \right) - d \left( \frac{m}{n} - \frac{1}{2} \right) + d \left( \frac{m}{n+1} - \frac{1}{2} \right) - d \left( \frac{m}{n} - \frac{1}{2} \right) \right) \geq 0 \quad (8)$$

In the latter case ( $\frac{m+1}{n} = \frac{1}{2}$ ) the mother compares her expected utility from having an additional child:

$$u_{c=1}^2 = f(n+1) + \alpha g \left( \frac{y}{p_i(n+1)} \right) - \frac{1}{2} \theta d \left( \frac{m+2}{n+1} - \frac{1}{2} \right) - \frac{1}{2} \theta d \left( \frac{m+1}{n+1} - \frac{1}{2} \right)$$

To her utility from not having an additional child:

$$u_{c=0}^2 = f(n) + \alpha g\left(\frac{y}{p;n}\right) - \theta d\left(\frac{m+1}{n} - \frac{1}{2}\right)$$

In other words, the mother will elect to continue fertility iff:

$$u_{c=1}^2 \geq u_{c=0}^2$$

$$\implies b_{diff}(n+1, n) - \frac{1}{2}\theta \left( d\left(\frac{m+2}{n+1} - \frac{1}{2}\right) - d\left(\frac{m+1}{n} - \frac{1}{2}\right) + d\left(\frac{m+1}{n+1} - \frac{1}{2}\right) - d\left(\frac{m+1}{n} - \frac{1}{2}\right) \right) \geq 0 \quad (9)$$

This simplification to  $b_{diff}(n+1, n)$  is possible because there are no direct impacts of the sex ratio on the resources allocated to each child or their health and because in both cases the total number of children is the same. Next, observe that by the concavity of  $d(\cdot)$  and making use of the fact that  $\frac{m+1}{n} = \frac{1}{2}$ :

$$\begin{aligned} & d\left(\frac{m+1}{n+1} - \frac{1}{2}\right) - d\left(\frac{m}{n} - \frac{1}{2}\right) + d\left(\frac{m}{n+1} - \frac{1}{2}\right) - d\left(\frac{m}{n} - \frac{1}{2}\right) \leq \\ & d\left(\frac{m+2}{n+1} - \frac{1}{2}\right) - d\left(\frac{m+1}{n} - \frac{1}{2}\right) + d\left(\frac{m+1}{n+1} - \frac{1}{2}\right) - d\left(\frac{m+1}{n} - \frac{1}{2}\right) \end{aligned}$$

Together, this implies that if:

$$u_{c=1}^2 \geq u_{c=0}^2 \equiv b_{diff}(n+1, n) - \frac{1}{2}\theta \left( d\left(\frac{m+2}{n+1} - \frac{1}{2}\right) - d\left(\frac{m+1}{n} - \frac{1}{2}\right) + d\left(\frac{m+1}{n+1} - \frac{1}{2}\right) - d\left(\frac{m+1}{n} - \frac{1}{2}\right) \right) \geq 0$$

Then it must also be true that:

$$u_{c=1}^1 \geq u_{c=0}^1 \equiv b_{diff}(n+1, n) - \frac{1}{2}\theta \left( d\left(\frac{m+1}{n+1} - \frac{1}{2}\right) - d\left(\frac{m}{n} - \frac{1}{2}\right) + d\left(\frac{m}{n+1} - \frac{1}{2}\right) - d\left(\frac{m}{n} - \frac{1}{2}\right) \right) \geq 0$$

So that any mother who would continue her fertility with a balanced existing sex ratio would also continue her fertility with a sex ratio that is lower than her ideal. I omit the identical proof for the symmetric statement that any mother who would continue her fertility with a balanced sex ratio would also continue her fertility with a sex ratio that is higher than ideal. Together these imply that mothers with balanced sex ratios are at most equally likely to continue fertility as mothers that are not at their ideal sex ratio.

Clearly, the converse statements need not be true. A mother who would continue fertility with an unbalanced sex ratio would not necessarily also continue fertility with

a balanced sex ratio. This is the case when:

$$u_{c=1}^1 \geq u_{c=0}^1 \text{ but } u_{c=1}^2 < u_{c=0}^2 \quad (10)$$

Which is not ruled out on the basis of assumptions (4)-(7).

**Proposition 4** If the increase in marginal utility from an additional child is low enough, mothers will elect not to continue fertility regardless of whether they have a balanced existing sex ratio.

Though this model assumes away direct utility or financial costs of additional children, each additional child imposes an indirect cost through the decrease in available resources per child resulting from the increased family size. Thus, a low increase in marginal utility implies that  $b_{diff}(n+1, n) \equiv f(n+1) + \alpha g\left(\frac{y}{p_i(n+1)}\right) - f(n) - \alpha g\left(\frac{y}{p_i n}\right)$  is small, or even negative. A sufficiently low value will imply that mothers elect not to continue fertility even without attaining their ideal sex ratio.

At what parity is  $b_{diff}(n+1, n)$  likely to be low? By (6), the net benefit of an additional child is decreasing in the existing number of children. This implies that the value of  $b_{diff}(n+1, n)$  will be low at higher parities. Intuitively, at high parities mothers, receiving less marginal benefit from a new child, are likely to end their fertility regardless of the sex composition among existing children. The expected benefit from potentially moving closer to their ideal sex ratio must be exceptionally large to compensate for the smaller increase in utility resulting from the additional child. Therefore, fertility will only differ with the mother's distance from her ideal sex ratio when the marginal benefit from a new child is sufficiently high, i.e. at lower birth parities.

**Proposition 5** Children who are born of their mother's preferred sex and at low birth parities will be healthier (higher quality) than children who are born of their mother's less preferred sex at low birth parities.

From above, mothers stated sex preference will reflect the sex composition among existing children. Those children who are born of their mother's preferred sex will move their mothers closer to the ideal sex composition,  $\gamma$ . Therefore, following the birth of this child, the mother who had a child of her preferred sex will be less likely to continue fertility and have another child as long as the set of mothers for which (10) is true is non-empty. From **Proposition 4**, this set of mothers is less likely to be empty at low birth parities. Of the children whose mothers satisfy (10), those who were born of their mother's preferred sex will split household resources  $y$  with at least one fewer child resulting in healthier (higher quality) children on average:  $\bar{q}_n \geq \bar{q}_{n+1}$ .

### **The Burden of Continued Fertility**

The model above predicts that all children in households where the most recent child was born of the mother's preferred sex would be, on average, of worse health or lower quality. It seems reasonable, however, to expect that a disproportionate amount of the difference would be borne by the most recent child instead of older siblings. This is because, while the model abstracts away from differences in the production function of health/quality, the economics literature largely agrees that health and skills are most sensitive during early childhood.<sup>23</sup> Therefore, if the introduction of a new child results in an equal reduction in resources for all existing children, it would likely still disproportionately impact the health of the youngest existing child. In fact, as long as mothers do not completely account for the gradient in health/quality productivity by age,<sup>24</sup> the youngest child should be the most impacted by continued fertility.

---

<sup>23</sup>See the main text of this paper for a more detailed discussion.

<sup>24</sup>This would be the case if mothers value, at least somewhat, ex-ante equality (i.e. they would prefer to treat all children equally), or if the inputs for older children are less substitutable with the inputs for younger children (i.e. solid food is less useful for infants than for teenagers).

# 13 Appendix B: Additional Tables

**TABLE B1: Fully Saturated Controls for Existing Sex Composition of Older Children**

	BMI Z-Score (1)	Weight-for-age Z-Score (2)	General Health (3)	Days of Primary Activity Missed (4)	Child is Wasted (5)
<u><i>Without Controls</i></u>					
Born of Mother's Preferred Sex	0.236*** (0.086)	0.164** (0.075)	-0.091* (0.052)	-0.499*** (0.170)	-0.032** (0.016)
<u><i>With Saturated HH Composition Controls</i></u>					
Born of Mother's Preferred Sex	0.256* (0.133)	0.127 (0.119)	-0.094 (0.091)	-0.596** (0.291)	-0.023 (0.026)
Observations	1,414	1,430	1,433	1,432	1,414
P-Value: Coefficients Are the Same	0.853	0.708	0.970	0.689	0.696

Notes: The sample includes children conceived immediately following the elicitation of their mothers preferences over subsequent child gender. In addition, their mothers must have either preferred male children or female children, they could not have been indifferent to either sex. Both panels present the results of OLS regressions. The without controls panel includes only the "Born of Mother's Preferred Sex" indicator and a constant. The with controls panel includes a full set of indicators for the number of older brothers and the number of older sisters each child has, and a full set of interactions. In addition, each column shows the P-value from a test of whether the coefficient on "Born of Mother's Preferred Sex" is equal in the two panels. Robust standard errors clustered at the mother level are in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.



**TABLE B2:**  
**Effect Heterogeneity: By Whether Child is Above or Below Birth Parity Three**

	BMI Z-Score (1)	Weight-for-Age Z-Score (2)
Above Parity Three	-0.240 (0.215)	-0.116 (0.183)
Parity Three and Below	0.318*** (0.094)	0.213*** (0.082)
Observations	1,414	1,414
P-value Individual Test Effects Are Equal	0.0170	0.101
P-value Joint Test That Effects Are Equal		0.0530

Notes: The sample includes children conceived immediately following the elicitation of their mothers' preferences over subsequent child gender. In addition, their mothers must have either preferred male children or female children, they could not have been indifferent to either sex. The first row of results estimates the effect of being born of the mother's preferred sex for children born at parities four and above. The second row of estimates does the same for children born at parities one through three. Both estimates are the result of OLS regressions. In addition, each column shows the P-value from a test of whether the coefficient on "Born of Mother's Preferred Sex" is equal for children who are born at parities one through three and children who are born at parities four and above. Finally, at the bottom of the table is the P-value from the joint test of whether both estimates are equal for children up to parity three and those above parity three. Robust standard errors clustered at the mother level are in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.


**Table B3: Older Sibling Health and Maternal Sex Preferences**

	BMI-for-Age Z-Score (1)	Weight-for-Age Z-Score (2)	Height-for-Age Z-Score (3)
<i><u>Mother Desires a Daughter</u></i>			
Average Anthropometric Z-Score For Older Sisters	-0.00826 (0.013)	-0.00428 (0.013)	-0.000977 (0.012)
Average Anthropometric Z-Score For Older Brothers	0.00739 (0.012)	-0.0184 (0.014)	-0.0233** (0.010)
<i><u>Mother Desires a Son</u></i>			
Average Anthropometric Z-Score For Older Sisters	0.0177 (0.011)	0.00143 (0.009)	-0.0115 (0.009)
Average Anthropometric Z-Score For Older Brothers	-0.00878 (0.009)	0.00597 (0.013)	0.0152 (0.012)
Observations	593	593	593

Notes: The sample includes children conceived immediately following the elicitation of their mothers preferences over subsequent child gender. In addition, their mothers must have either preferred male children or female children, they could not have been indifferent to either sex, and children must have at least one older brother and one older sister with non-missing anthropometric measurements before the child's birth. Both panels present the results of Linear Probability Models. The sibling anthropometric measure compared is listed at the top of each column. The top panel compares average older sibling anthropometric measures for mothers who desired a girl next; the bottom panel for mothers who desired a boy next. Robust standard errors clustered at the mother level are in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.

## 14 Appendix C: Additional Figures

Figure C1: Maternal Sex Preference and Birth Spacing



Note: Figure C1 Panel A presents the proportion of the main sample children whose mothers have had an additional child at the given number of months after the birth of the main child. The number of months since the main child's birth is shown on the horizontal axis. The Kaplan-Meier cumulative failure (additional birth) functions are plotted separately by whether the main child was born of the mother's (pre-conception) preferred or less preferred sex. Duration is top coded to 72 months. Figure C1 Panel B plots the same Kaplan-Meier cumulative failure function but limiting the sample to mothers who have at least one additional child. Panel B represents purely the intensive birth spacing margin.

## 15 Appendix D: Identification of Treatment Effects

Let  $T_i \in \{0, 1\}$  represent whether child  $i$  is born of their mother's pre-conception preferred sex. Denote whether child  $i$  is male using the male indicator  $M_i \in \{0, 1\}$  and whether child  $i$ 's mother preferred to have a male child before conception using  $PM_i \in \{0, 1\}$ .

$$T_i = M_i PM_i + (1 - M_i)(1 - PM_i)$$

Define the potential early health outcomes for child  $i$  as:

$$y_{i1} = y(T_i = 1), y_{i0} = y(T_i = 0)$$

Finally, I observe:

$$y_i = y_{i1} T_i + y_{i0} (1 - T_i)$$

I would like to estimate the average treatment effect of being born of the mother's pre-conception preferred sex:

$$\tau_{ATE} = E[y_1 - y_0]$$

Empirically, I observe:

$$\hat{\tau} = E[y_i | T_i = 1] - E[y_i | T_i = 0]$$

I make the following assumptions:

$$Pr\{PM_i = 1\} = Pr\{PM_i = 0\} = \frac{1}{2} \tag{11}$$

$$Pr\{M_i = 1 | PM_i = 1\} = Pr\{M_i = 0 | PM_i = 0\} \tag{12}$$

$$y_1, y_0 \perp M_i | PM_i \quad (13)$$

Assumption (11) requires that there not be any son (or daughter) preference in the sample. Table 1 Panel A shows that this assumption is met. Of mothers who have a gender preference, fifty percent of mothers in the sample desire a female child next and the other fifty percent desire a male child. Under Assumption (12), the likelihood of having a child of the preferred sex must be the same for mothers who desire a son next and mothers who desire a daughter next. The empirical justification for Assumption (12) can be found in Table 2.  $Pr\{M_i = 1 | PM_i = 1\}$ , the probability that the next child is male given that the mother desired a male is .47;  $Pr\{M_i = 0 | PM_i = 0\}$ , the probability that the next child is female given that the mother desired a female is  $1 - .533 = .467$ . These probabilities are nearly identical and are not significantly different from one another. Also implied by Assumption (12) and the information contained in Table 2 is that  $Pr\{M_i = 1 | PM_i = 0\} = Pr\{M_i = 0 | PM_i = 1\}$ .

Assumption (13), the conditional independence assumption in this context, is necessary because identification uses the random allocation of gender at conception to create variation in the treatment. In practice, I need the slightly weaker assumption of mean independence of potential outcomes and treatment conditional on the mother's preferred sex. As discussed in the main text, there is no evidence of sex-selective abortion in the sample. As a result, any violation of Assumption (13) would require mothers to form accurate expectations about the relative health of not-yet-conceived male and female children and for these expectations to be reflected in mothers' stated sex preferences; there are few plausible scenarios that satisfy these requirements.

One candidate, given the evidence that stated sex preferences are largely determined by the existing sex distribution in the household, is that there are direct spillovers from older siblings that impact the early health of the next child. If, for example, having more older sisters negatively impacts the early health of subsequent female children for reasons unrelated to parental behavior, then the estimated treatment effects without any additional

controls will be biased. However, as Table B1 in Appendix B demonstrates, direct effects of the previous sex composition are not driving the estimated treatment effects; including a fully saturated control for the existing sibling composition has no effect on the estimates. Therefore, it is unlikely that spillovers from the older sibling composition violate (13).

Alternatively, assumption (13) would be violated if mothers desire a particular gender for their next child because they expect to produce healthier children of that sex for genetic reasons. If parents believe they have a genetic advantage at producing a particular sex, their beliefs about this comparative advantage are correct, and it makes them more likely to desire a child of that sex, then my estimates will wrongly attribute this correlation between stated sex preference and sex-specific genetic advantage to parental behavior.

There are a number of reasons to be skeptical of this possibility. First, as discussed above, stated sex preferences are largely determined by the existing sex distribution in the household; observed fertility and stated sex preferences are consistent with parents reacting to sex realizations in an attempt to achieve a balanced sex ratio. Second, if parents are correct about having a genetic advantage at producing children of one gender, I would expect to see that advantage reflected in child endowments at birth. Specifically, I should observe that children born of their mother's preferred sex are heavier at birth. Table 6 shows that children that are born of their mother's preferred sex have no advantage in birth weight. The treatment effects must therefore be the result of differential growth after birth. Finally, if mothers form expectations about their genetic advantage at producing males or females, I would expect the health of similarly-gendered older siblings to be an informative signal. That is, mothers who have relatively healthier existing daughters than sons should be more likely to demand a female child next.

To test whether this is the case, Appendix Table B3 checks whether the body mass index (BMI), weight, and height of older siblings, measured prior to the conception of the main child, have any effect on the likelihood that the mother preferred for the main child to be male (or female). In each column, the sample is limited to mothers who have

at least one older male child and one older female child with non-missing anthropometric measures. BMI, weight, and height are standardized by sex and age in months using the same WHO growth standards from 2006/2007 that were used to standardize outcomes for the children used in the main analysis. The Table presents the estimated effects of the anthropometric measure listed at the top of the column on the probability the mother wants a male (or female) child next. All of the twelve coefficients are small in magnitude and only one is even marginally statistically significant: mothers with taller existing male children appear slightly less likely to desire a daughter next. However, mothers with taller sons are not more likely to desire a son next. Overall, there does not appear to be evidence that mothers with healthier existing daughters are more likely to desire future daughters, nor that mothers with healthier existing sons are more likely to desire future sons as a violation of Assumption (13) would predict.

Therefore, the available evidence fails to support any of the likely violations to Assumption (13). Making use of Assumptions (11), (12), and (13), the proof that  $\hat{\tau} = \tau_{ATE}$  is presented below.

$$\hat{\tau} = E[y_i|T_i = 1] - E[y_i|T_i = 0]$$

Focus on the first term:

$$\begin{aligned} E[y_i|T_i = 1] &= E[y_1|T_i = 1] = E[y_1|M_i = 1, PM_i = 1] \frac{Pr\{M_i = 1, PM_i = 1\}}{Pr\{M_i = 1, PM_i = 1\} + Pr\{M_i = 0, PM_i = 0\}} \\ &+ E[y_1|M_i = 0, PM_i = 0] \frac{Pr\{M_i = 0, PM_i = 0\}}{Pr\{M_i = 1, PM_i = 1\} + Pr\{M_i = 0, PM_i = 0\}} \end{aligned}$$

Using (13) this simplifies to:

$$E[y_1|PM_i = 1] \frac{Pr\{M_i = 1, PM_i = 1\}}{Pr\{M_i = 1, PM_i = 1\} + Pr\{M_i = 0, PM_i = 0\}}$$

$$+E[y_1|PM_i = 0] \frac{Pr\{M_i = 0, PM_i = 0\}}{Pr\{M_i = 1, PM_i = 1\} + Pr\{M_i = 0, PM_i = 0\}}$$

Using (12) this further simplifies to:

$$E[y_1|PM_i = 1] \frac{1}{2} + E[y_1|PM_i = 0] \frac{1}{2}$$

Now use (11) to rewrite the above as:

$$E[y_1|PM_i = 1] Pr\{PM_i = 1\} + E[y_1|PM_i = 0] Pr\{PM_i = 0\}$$

Using basic properties of expectations this is equivalent to:

$$E[y_1|PM_i = 1] Pr\{PM_i = 1\} + E[y_1|PM_i = 0] Pr\{PM_i = 0\} = E[y_1]$$

Following analogous steps yields the following simplification for the second term:

$$E[y_0|T_i = 0] = E[y_0|M_i = 1, PM_i = 0] Pr\{PM_i = 0\} + E[y_0|M_i = 0, PM_i = 1] Pr\{PM_i = 1\}$$

$$= E[y_0|PM_i = 0] Pr\{PM_i = 0\} + E[y_0|PM_i = 1] Pr\{PM_i = 1\}$$

$$= E[y_0]$$

Combining the above steps implies that:

$$\hat{\tau} = E[y_i|T_i = 1] - E[y_i|T_i = 0] = E[y_1] - E[y_0] = \tau_{ATE}$$

$$\implies \hat{\tau} = \tau_{ATE}$$



## Part II

# Parental Response to In Utero Shocks

## 1 Introduction

Until recently, unresolved empirical issues and a lack of adequate data prevented researchers from claiming that correlations between poor early childhood conditions and poor adult outcomes represented causal effects. In particular, researchers struggled to identify whether any component of adult outcomes was caused by early childhood conditions or if both were simply influenced by the same unobserved factors.<sup>25</sup>

Recent developments in econometrics and an increased availability of large, detailed panel data sets have made the identification of causal effects more feasible. Research across a variety of disciplines now shows that poor early childhood conditions cause worse adult health [Barker and Clark, 1997, Currie and Hyson, 1999, Almond and Mazumder, 2005, Case et al., 2005] lower educational attainment and cognitive ability [Almond et al., 2009, Maccini and Yang, 2009, Walker et al., 2011, Akresh et al., 2012], and depressed labor market earnings [Black et al., 2007, Cunha and Heckman, 2007, Cunha et al., 2010, Gertler et al., 2013].<sup>26</sup> Clearly, early conditions<sup>27</sup> and exogenous shocks in early childhood<sup>28</sup> affect multiple dimensions of well-being in adulthood. From the perspective of intergenerational economic mobility, an individual's social, economic, and health outcomes are limited by factors beyond her control.

This paper contributes to the literature by examining the relationship between one type of exogenous shock in early childhood, rainfall exposure during the in utero period, and

---

<sup>25</sup>For example, persistent negative health shocks, low values of parental inputs, or unobserved genetic components.

<sup>26</sup>See Almond and Currie [2010] for a review of the literature.

<sup>27</sup>e.g. schooling, early nutrition, and the general home environment.

<sup>28</sup>e.g. disasters, beneficial program rollouts, and random variation in weather conditions.

long-term outcomes (height, weight, schooling attainment, and cognitive ability). In contrast to prior work, which generally focuses on identifying long-term causal effects, I concentrate on estimating short-term behavioral responses by parents. Conditional on early parental behavior being important for long-term outcomes, behavioral changes in response to early shocks are evidence of an indirect causal channel running from the initial shock to parental behavior to long-term outcomes.

Most research on the long-term effects of early childhood shocks focuses on developed countries. However, there are a number of reasons to think that these effects should be even stronger in developing countries [Currie and Vogl, 2013]. Children in developing countries are both more likely to experience negative shocks during childhood and less likely to live in households that have the resources to remediate the initial impact of these shocks. Within the development economics literature, several papers have used the exogenous timing of the expansion of government programs to study the effects of early conditions on medium to long-term outcomes [Maluccio et al., 2009, Ozier, 2011, Macours et al., 2012]. In general, these studies find that exposure to health interventions<sup>29</sup> leads to improved cognitive ability.

Another related set of papers tests whether variation in rainfall during early childhood has long-lasting effects on well-being. In many developing countries, a majority of households rely on agriculture for their subsistence. As a result, rainfall can generate differences in child endowments that may persist into adulthood. For example, in a similar context to the one considered here, Maccini and Yang [2009] find that higher exposure to rainfall in Indonesia during the year after a child is born is associated with better health, school, and economic outcomes in adulthood. Interestingly, significant effects are found only for rural girls. Urban children, if anything, are negatively impacted by the early rainfall. The authors hypothesize that their finding is the result of a positive income effect (more rainfall is associated with higher crop yields) during a crucial period of child development. In combi-

---

<sup>29</sup>In each case the treatment is a health intervention. Macours et al. [2012] examine a nutritional supplementation program, Ozier [2011] studies a de-worming program, and Field et al. [2009] evaluate an iodine supplementation program.

nation with gender biased preferences among parents, less rainfall leads parents to reduce nutrition for girls but not for boys. Maccini and Yang [2009] suggest that a reduction in available nutrition may lead to worse early health among the rural girls in their sample. However, they do not explicitly test for changes in parental behavior that might lead to worse early health conditions.

In Burkina Faso, Akresh et al. [2012] use measures of exposure to drought while in utero as the treatment. They find that children exposed to drought in utero perform significantly worse on cognitive tests than their unaffected siblings later in life. They attribute this effect to a combination of epigenetic factors and a sibling rivalry for scarce household resources. As with all household fixed effects strategies, measurement error is potentially an important issue.

Using a different empirical methodology, Millett and Shah [2012a] find that children in rural India who are exposed to drought in utero perform worse than non-exposed children of similar ages on both math and reading tests later in childhood. A lengthy time series of rainfall data allows them to compare the importance of rainfall exposure in utero with rainfall exposure later in childhood. They conclude that rainfall exposure only matters in utero and during the first two years of life. A limitation of Millett and Shah's analysis is that the data do not contain precise information on date of birth. They are forced to assume that all children are born at the midpoint of each year. At best, if there are no unobserved seasonal differences in the "quality" of children, this strategy will lead to attenuated coefficients.<sup>30</sup> Worse, as the authors note, if season of birth is correlated with unobserved characteristics that are also determinants of ability, their results are biased in an ambiguous direction.<sup>31</sup>

In related work, Millett and Shah [2012b] test whether rainfall shocks might negatively affect child ability as measured by test scores. Increased rainfall creates both a positive income effect and a negative substitution effect. Specifically, parents may supply more labor

---

<sup>30</sup>The authors are aware of this possibility.

<sup>31</sup>See Buckles and Hungerman [2010] and Pitt and Sigle [2012] for two examples of how season of birth might be important.

in response to the increase in the marginal productivity of labor generated by rainfall. As a result they may spend less time on activities that build human capital in children, creating a negative substitution effect. The net effect of increased rainfall is therefore unclear. Using data from India, they find evidence to support the idea that higher contemporaneous rainfall reduces school attendance and measured cognitive ability among children. In contrast, I use data that contain month and year of birth. Using the precise month of birth, I generate month by month measures of in utero rainfall exposure for each child. Therefore, my results should not suffer from measurement error in the timing of birth. I explore empirically whether initial health is more influenced by the income effect or the substitution effect and direct, negative health effects that epidemiological studies suggest may result from an increase in utero rainfall.<sup>32</sup>

This paper contributes to the literature on reinforcing and compensating behavior by parents. Within economics, this literature can be traced back to the seminal work by Becker and Tomes [1976]. The authors provide a theoretical framework for how and why parents might adjust investment (both time and monetary) in response to child characteristics.

More recently, Almond and Currie [2010] provide a simple model to illustrate compensating and reinforcing behavior by parents. In the two period version of their model, human capital in period  $t$  is a function of parental investment in contemporaneous and previous periods, several fixed parameters, and the realization of stochastic shocks. They show that parents' decisions to reinforce<sup>33</sup> or compensate early shocks depend on the curvature of their utility function and the shape of the child's production function for human capital.

Several recent papers empirically test for compensating or reinforcing behavior by parents. Royer [2009] uses data from the Early Childhood Longitudinal Program, Birth Cohort (ECLS-B) containing information on 856 twins to test for differences in parental investments. She finds no significant differences in parental investment<sup>34</sup> behavior in response

---

<sup>32</sup>A third potential mechanism is worse maternal health. This channel is discussed in more detail below.

<sup>33</sup>That is, to respond to shocks by shifting investment in the same direction. E.g. respond to a positive shock by increasing investment in the child or respond to a negative shock by decreasing investment.

<sup>34</sup>Royer considers two dimensions of early medical care: NICU use and the number of days spent in the

to differences in birth weight between twins. Bharadwaj et al. [2012] use a regression discontinuity approach to test whether being classified as very low birth weight has any effect on parental investments. The authors find no evidence that parents of children just below the very low birth weight cutoff invest differently than parents of children just above the cutoff.

Finally, Adhvaryu and Nyshadham [2012] consider how exposing a child to an iodine supplementation program affects parental behavior in Tanzania. First, the authors show that the supplementation program had significant, positive effects on long-term outcomes (years of school completed) for their sample. Then, they test whether parental behavior is responsive to the change in early childhood conditions induced by program exposure. The authors find that exposure to the iodine program in utero significantly increased vaccination take-up and the length of breastfeeding. My paper compliments the literature on compensating an reinforcing behavior by exploring whether parental investment responds to exogenous negative shocks in early childhood. My analysis is distinct from that of Adhvaryu and Nyshadham [2012] in that I consider how parents respond to negative shocks as well as positive shocks. Parents might respond asymmetrically to shocks if, for example, they have preferences that are asymmetric around a reference point for child quality or if they suffer from credit constraints that are differentially affected by positive and negative shocks.

Nearly all of the previous work estimating behavioral responses by parents uses either extreme environmental events (e.g. droughts, epidemics, extreme birth outcomes)<sup>35</sup> or randomized interventions that are paired with informational treatments.<sup>36</sup> While interesting, extreme events are by definition rare. We might expect parents to respond quite differently in extreme circumstances than they would to more typical environmental variation. The rainfall data I use have the advantage of capturing more routine variation for households in

---

hospital.

<sup>35</sup>See Akresh et al. [2012], Bleakley [2010], Bharadwaj et al. [2012] for examples.

<sup>36</sup>See Adhvaryu and Nyshadham [2012], Aizer and Cunha [2012] for examples.

developing countries. From a policy perspective, parental responses to both types of variation are potentially important, but estimates currently only exist for how parents adjust to extreme events.

Randomized interventions provide clean identification of average treatment effect parameters. However, because they are generally paired with educational treatments that may also shift parental preferences and expectations, the identified treatment effects average parental responses to changes in early child quality (e.g. health, ability) and the interaction between changes in child quality and changes in parental preferences/expectations. In contrast, rainfall deviation is unlikely to alter parental preferences in a meaningful way. Thus, under modest assumptions, my results can be interpreted as capturing how parents respond to exogenous changes in early childhood health.

The remainder of the paper proceeds as follows. Section 2 discusses the importance of rainfall in Indonesia and, more generally, in developing countries. Section 3 provides a brief theoretical framework while Section 4 summarizes the data. Section 5 presents results and Section 6 concludes.

## 2 Rainfall

In the early 1990's, approximately half of the 180 million inhabitants of Indonesia were employed in agriculture. Though this figure has gradually decreased over the past two decades, the agricultural sector still represents 15% of Indonesian GDP[CIA, 2012], a majority of which (56.5%) comes from farms under a half hectare in size[Sudaryanto et al., 2009].<sup>37</sup> Importantly, these small farms are less likely to have access to productivity enhancing physical capital and less likely to be irrigated. Rainfall is thus crucial for agricultural production. Rice, the most common crop in Indonesia, is particularly sensitive to rainfall. Therefore, rice output in Indonesia, like rainfall, is highly seasonal and fluctuates

---

<sup>37</sup>These figures likely understate the true importance of agricultural output in Indonesia because of the relatively low value added of agricultural labor.

both across seasons and across geographic locations within the same season. Figure 1 plots the average rainfall by month over the 1970-1993 period for the province of West Java. I observe a clear wet season (November through April) and a clear dry season (May through October). The wet season months and dry season months vary slightly across provinces. Though not shown here, there is also likely to be variation in the timing of the rainy season within provinces. There is also variation within provinces across years. The variation in the timing of the wet and dry seasons causes important changes in farmer behavior. For example, farmers must wait until enough rain has accumulated to flood their fields before they can seed. Late arrival of the rainy season therefore shortens the growing season. Rice also requires a considerable amount of water to grow after seeding. Early arrival of the dry season will also shorten the growing season and reduce output.

Empirical evidence underscores the sensitivity of agriculture, and rice in particular, to fluctuations in rainfall. Parchure [2002] suggest that in India nearly 90% of agricultural output variation can be attributed to rainfall. In Indonesia, Levine and Yang [2006] find that 10% higher rainfall in a district is associated with an increase of 0.4% in the agricultural output of rice. Further analysis by Maccini and Yang [2009] shows that this positive relationship is monotonic. High levels of rainfall do not appear to hurt the production of rice in Indonesia. While not all Indonesian adults are directly involved in agricultural production, fluctuations in agricultural productivity have welfare implications for nearly all Indonesians. For adults involved in subsistence farming, better harvests improve welfare directly through an income effect. For adults who supply labor in agricultural markets, agricultural productivity appears to be even more important [Jayachandran, 2006]. Even adults who are not directly involved in agriculture may be indirectly affected through price effects. In years when harvests are plentiful food is cheaper (all else equal). This should shift budget sets outwards making net purchasers unambiguously better off. As I discuss below, the effect of rainfall fluctuations on child welfare is less clear.

Highly variant rainfall, across both time and space, coupled with the strong relationship

between rainfall and agricultural output, implies that crop production will vary considerably across provinces within a year and within provinces across time. If households do not have access to complete formal or informal risk-sharing instruments, consumption will also fluctuate over time. A large literature in development economics focuses on whether households are able to perfectly smooth consumption and income.<sup>38</sup> The consensus in the literature is that risk-sharing is incomplete for most households in developing countries. As a result, there is strong reason to believe that deviations in rainfall will have important effects on the level of nutrition available in households.

### **3 Theoretical Framework**

#### **The Initial Effect of Rainfall Shocks**

Rainfall exposure in utero has a theoretically ambiguous effect on early childhood health.<sup>39</sup> On the one hand, rainfall should increase the productivity of labor in agriculture. Agricultural laborers should see their wages increase<sup>40</sup> and farm owners (subsistence and otherwise) should see their harvest and profits increase. If households are unable to perfectly smooth consumption across periods, this results in a positive income effect on fetal health through an increase in available nutrition. Research in public health, economics, and medicine shows that lack of nutrition during the in utero period is associated with lower initial health endowments [Kowaleski-Jones and Duncan, 2002, Bitler and Currie, 2005a,b, Figlio et al., 2009, Almond et al., 2011, Almond and Mazumder, 2011, Hoynes et al., 2011].<sup>41</sup>

---

<sup>38</sup>See Mazzocco and Saini, 2012, Morduch, 1995, Ravallion and Chaudhuri, 1997, Rosenzweig, 1988, Rosenzweig and Stark, 1989, Rosenzweig and Wolpin, 1993, Townsend, 1994 for examples.

<sup>39</sup>I use the terms early childhood health and initial endowments interchangeably. In both cases I am referring to the three measures discussed above: birth weight, relative size at birth, and the index combining birth weight and relative size.

<sup>40</sup>However, there is some evidence that agricultural wages are asymmetrically sticky [Jayachandran, 2006].

<sup>41</sup>In general, the literature uses birth weight, APGAR scores, and gestational length as measures of initial health.



However, as pointed out by Maccini and Yang [2009] and Millett and Shah [2012b], increased rainfall also generates a negative substitution effect by raising the opportunity cost of parental time. The increase in agricultural productivity increases the marginal return to supplying labor. This increases parental labor supply and, potentially, reduces the amount of time spent on health seeking behaviors. If parental time inputs, and particularly maternal time input, are important during the in utero period then an increase in rainfall may negatively affect initial child health. As a result of how crops are grown in Indonesia, labor supply is most likely to increase immediately following extended rainfall (when farmers are seeding) and three months later (when it is time to harvest).

A third possibility is that an increase in rainfall increases the likelihood of illness for the mother [Hunter, 1997, Curriero et al., 2001]. Maternal infection can divert energy and nutrition away from the womb and towards fighting off the infection. This is harmful to development of the fetus. Ex-ante it is unclear which of these three effects will dominate but there are some testable predictions that arise from the theory.

1. If initial health are positively affected by in utero rainfall exposure the income effect dominates.
2. If initial health are negatively affected by in utero rainfall exposure it must be a result of either the negative substitution effect or the increased likelihood of disease.
3. If the results are negative *and* more pronounced among households that are a priori more susceptible to waterborne illnesses then the increased probability of disease is likely to be the driving factor.

The first two predictions are straightforward. If the income effect dominates we should expect to see initial health endowments improved by in utero rainfall exposure. If not, higher rainfall in utero should be associated with worse initial endowments.

The third prediction is less obvious. I do not currently have reliable measures of maternal health during pregnancy. If I did, I would test whether maternal health during pregnancy is negatively affected by rainfall exposure. However, if increased rainfall exposure in utero

negatively affects the measures of initial endowments, I can test whether this effect is more pronounced among households likely to be at high risk for waterborne diseases. If the negative impact of rainfall is driven by high risk households, then a positive association between rainfall and the probability of the mother contracting an illness while pregnant is the most likely mechanism.

Further complicating the analysis is the possibility that in utero rainfall may have a different effect depending on gestational age at the time of exposure. In general, in developing countries there appears to be a strong positive relationship between maternal nutrition and birth outcomes.<sup>42</sup> However, within the nine months prior to birth, maternal pre-natal nutrition may be relatively more important during certain months and maternal rest during others. In practice, I use separate rainfall measures for each trimester of pregnancy and therefore empirically test whether there are heterogeneous effects by gestational age.

### **The Long-Term Effect of Rainfall Shocks**

The long-term effect of rainfall shocks depends on both the initial effect on infant health and how parental investment after birth responds to the variation in early child health caused by the shocks. One possibility is that there is no direct, long-term effect of shocks in early childhood. Moreover, parental investment post-birth may or may not respond to the changes in health induced by early shocks. If parental investment does respond to changes in early health, it is also unclear in which direction it will respond. Theory suggests that parents' responses depend on both the shape of the production function for child quality<sup>43</sup> and the shape of the parents' utility function. Specifically, the response depends on the parents' perception of how effective investments in early childhood are at remediating lower initial endowments. The more effective parents perceive their investments to be the more likely it is that parents will compensate for initial negative shocks. With standard assumptions on

---

<sup>42</sup>See Abu-Saad and Fraser, 2010, Ceesay et al., 1997 for evidence from the epidemiological and medical literature.

<sup>43</sup>In the context of this paper, I assume child quality can be represented by health and more general human capital.

parental utility and some substitutability between initial conditions and early investment, negative shocks to early conditions should cause parents to increase early investment. By reducing initial child quality, negative shocks increase the marginal utility of child quality relative to consumption. Thus compensating behavior should be expected [Almond and Currie, 2010]. On the other hand, if parents are credit constrained, or if there are multiple children in the household, parents might reinforce initial shocks instead.<sup>44</sup>

To help illustrate what contributes to compensating and reinforcing behavior, I present a simple, static, deterministic model of the post-birth parental investment decision. I assume that parents receive utility from child quality and from their consumption of other goods. The parents' maximization problem is to choose the level of investment in child quality and consumption of other goods subject to the budget constraint and exogenous income,  $y$ :

$$\max_{i_s, c} u(q_s, c) \text{ s.t. } y \geq pi_s + c, q_s = f(i_s, \varepsilon_s) \quad (14)$$

Where  $q_s, s \in \{f, m\}$ , represents the quality of a child of sex  $s$ , and  $c$  is a numeraire good.  $q_s$  is determined by  $q_s = f(i_s, \varepsilon_s)$  where  $\varepsilon_s$  captures the newborn's initial health which is modeled as an exogenous shock, and  $i_s$  is a measure of parental investment. Assuming Inada conditions on the utility and production functions,<sup>45</sup> first order conditions imply that at the solution:

$i_s$  :

$$u_q(f(i_s^*, \varepsilon_s), c^*) f_{i_s}(i_s^*, \varepsilon_s) - \lambda p = 0 \quad (15)$$

$c$  :

$$u_c(f(i_s^*, \varepsilon_s), c^*) - \lambda = 0 \quad (16)$$

---

<sup>44</sup>For example, if parents get significantly more utility from child quality once it reaches a certain threshold (being a college graduate, etc...), then reinforcing behavior may be more likely.

<sup>45</sup>Specifically,  $\lim_{x \rightarrow 0} u_c(q_s, x) = +\infty$ ,  $\lim_{z \rightarrow 0} u_q(z, c) = +\infty$ ,  $u_q > 0$ ,  $u_c > 0$ ,  $u_{qq} < 0$ ,  $u_{cc} < 0$ ,  $f(0, \varepsilon) = f(i, 0) = 0$ ,  $f_{ii} < 0$ ,  $f_{\varepsilon\varepsilon} < 0$ ,  $f_i > 0$ ,  $f_\varepsilon > 0$ .

$\lambda$  :

$$y = p i_s^* - c^* \quad (17)$$

Combining the first two conditions yields:

$$p u_c (f (i_s^*, \varepsilon_s), c^*) = u_q (f (i_s^*, \varepsilon_s), c^*) f_{i_s} (i_s^*, \varepsilon_s) \quad (18)$$

Which offers the usual intuition. Parents invest in child quality until the increase in utility from an extra unit of investment is equated to the price scaled increase in utility from an extra unit of consumption of the numeraire good. In other words, parents consume until the marginal rate of substitution between the two goods equals the price ratio. For my purposes, I am interested in how parents adjust their optimal investment in response to changes in initial child quality. That is, I care about signing  $\frac{\partial i_s^*(\varepsilon)}{\partial \varepsilon}$ . To do so, I rewrite the two remaining first order conditions as functions of the choice variables and the child's initial health, and apply Cramer's rule to the resulting system. See Appendix A for details. After applying some algebra, I get the following expression for how parents adjust investment in response to shocks to initial child health:

$$\frac{\partial i_s^*(\varepsilon)}{\partial \varepsilon_s} = \frac{[p u_{cq} (f(\cdot), c^*(\cdot)) f_{\varepsilon_s}(\cdot) + u_{qq} (f(\cdot), c^*(\cdot)) f_{\varepsilon_s}(\cdot) f_{i_s}(\cdot) + u_q (f(\cdot), c^*(\cdot)) f_{i_s \varepsilon_s}(\cdot)]}{p [p u_{cc} (f(\cdot), c^*(\cdot)) - u_{qc} (f(\cdot), c^*(\cdot)) f_{i_s}(\cdot)] - [p u_{cq} (f(\cdot), c^*(\cdot)) f_{i_s}^*(\cdot) + u_{qq} (f(\cdot), c^*(\cdot)) f_{i_s}(\cdot)^2 + u_q (f(\cdot), c^*(\cdot)) f_{i_s i_s}(\cdot)]} \quad (19)$$

Without making any additional assumptions on either the shape of the parental utility function or the shape of the production function it is not possible to sign the impact that child quality shocks have on parental investment. Assuming that parental utility is additively separable in the numeraire good and child quality simplifies the expression considerably. The modified expression is given by:

$$\frac{\partial i_s^*(\varepsilon_s)}{\partial \varepsilon_s} = \frac{[u_{qq} (f(\cdot), c^*(\cdot)) f_{\varepsilon_s}(\cdot) f_{i_s}(\cdot) + u_q (f(\cdot), c^*(\cdot)) f_{i_s \varepsilon_s}(\cdot)]}{p [p u_{cc} (f(\cdot), c^*(\cdot))] - u_{qq} (f(\cdot), c^*(\cdot)) f_{i_s}(\cdot)^2 - u_q (f(\cdot), c^*(\cdot)) f_{i_s i_s}(\cdot)} \quad (20)$$

Based only on assumptions already made, this simplified expression still can not be signed. The first term in the numerator is unambiguously negative but the sign of the second term depends on whether initial endowments and parental investment are complements or substitutes in the production function of child quality. The sign of the denominator is also ambiguous and depends on the curvature of the utility function with respect to both arguments at the optimum and the curvature of the production function.

If I additionally assume that the production function of child quality is linear, the numerator is negative:

$$u_{qq}(f(\cdot), c^*(\cdot)) f_{\varepsilon_s}(\cdot) f_{i_s}(\cdot) + u_q(f(\cdot), c^*(\cdot)) f_{i_s \varepsilon_s}(\cdot) = u_{qq}(f(\cdot), c^*(\cdot)) f_{\varepsilon_s}(\cdot) f_{i_s}(\cdot) < 0 \quad (21)$$

Because the numerator is negative, the sign of the entire expression is the opposite of the sign of the denominator. That is:

$$\begin{aligned} \text{SIGN} \left\{ \frac{\partial i^*(\varepsilon)}{\partial \varepsilon} \right\} &= -\text{SIGN} \left\{ p [p u_{cc}(f(\cdot), c^*(\cdot))] - u_{qq}(f(\cdot), c^*(\cdot)) f_{i_s}(\cdot)^2 \right\} \quad (22) \\ &= \text{SIGN} \left\{ -p^2 u_{cc}(f(\cdot), c^*(\cdot)) + u_{qq}(f(\cdot), c^*(\cdot)) f_{i_s}(\cdot)^2 \right\} \end{aligned}$$

While still of ambiguous sign, there are several reasonable predictions that fall out of the model:

1. The likelihood that  $\frac{\partial i^*(\varepsilon)}{\partial \varepsilon} > 0$  is increasing in the price of investment.

This condition is quite intuitive: the more expensive it is for parents to remediate initial deficits the less likely it is that they engage in compensating behavior. This also offers a testable prediction for the model. In particular, as parental time becomes more valuable in the labor market do parents reduce the time investment they make in child quality? In the context of this paper where the majority of study subjects are involved in agricultural

production, I check whether parents reduce investment when the labor market returns to agriculture are likely to be high. Specifically, increased rainfall after initial child conditions are observed by the parents should increase labor market productivity and increase the opportunity cost of investment in children. The model predicts that more rainfall after birth should therefore increase the likelihood that parents reinforce the initial health conditions of the child. This can be estimated as the interaction between initial child health and a measure of post-birth rainfall in a model where the dependent variable is a measure of parental investment. A positive coefficient on the interaction would confirm the model's prediction.

2. The likelihood that  $\frac{\partial i^*(\varepsilon)}{\partial \varepsilon} > 0$  is increasing in  $ABS\{u_{cc}(f(\cdot), c^*(\cdot))\}$  and decreasing in  $ABS\{u_{qq}(f(\cdot), c^*(\cdot))\}$ .

The greater the curvature of the parental utility function with respect to consumption and the less the curvature of the utility function with respect to child quality, the more likely it is that parents reinforce initial conditions. In other words, if the utility cost of purchasing more child quality is low (relative to other consumption), taking into account both the direct cost and the indirect cost which reduces the marginal utility of consumption from additional consumption of child quality, parents are more likely to remediate initial shocks to child quality.

3. The likelihood that  $\frac{\partial i^*(\varepsilon)}{\partial \varepsilon} > 0$  is decreasing in  $f_{i_s}(\cdot)$ .

The greater the reduction in child quality resulting from a decrease in investment, the less likely it is that parents reduce their investment in response to low initial quality.

Clearly, even after making strong assumptions, theory does not provide an unambiguous prediction for how parents will respond to variation in initial child health. In practice, different parents will respond differently and the same parents may even respond differently to distinct children. With respect to the latter possibility, we might expect parents to respond differently to children that are observably different. One dimension of particular relevance is gender. Parents could respond differently to initial health for children of dif-

ferent genders if they perceive that the gradient of the production function for child quality is larger for one gender or if they value child quality more for one gender than another. I test whether parental responses differ by gender for the children in my sample. While this paper does not explicitly test for why parents compensate or reinforce initial child health, future work using similar data will focus on identifying the importance of elicited parental gender preferences for compensating and reinforcing behavior.

It is likely that the full effect of in utero shocks is not captured by my measures of initial child health. Health insults in utero may remain dormant for years before becoming observable. For example, the incidence of heart defects and type 2 diabetes in middle age have been linked to negative shocks in utero [Barker and Clark, 1997, Barker, 1998]. It seems plausible that there could be similar latent effects on cognitive ability or socio-emotional skills. I am not able to completely rule out this possibility. That said, no previous study has found evidence of such latent effects. In any case, parents will not observe these latent effects and therefore will not respond by shifting investment.

## **4 Data**

### **Rainfall Data**

Obtaining detailed rainfall data is of critical importance to this project. The data I use are from rainfall stations operated by the Indonesian Meteorological and Geophysical Agency. The raw data were collected and organized with cooperation from the Japan Agency for Marine-Earth Science and Technology. I begin with measures of daily rainfall in millimeters from 157 rainfall stations across Indonesia. For some stations, I have data covering all of 1961-1993. Unfortunately, most rainfall stations are missing some data, most often during the earlier years. In response, I drop all years before 1970.<sup>46</sup> I generate monthly rainfall

---

<sup>46</sup>The data are also somewhat sparse for the years between 1991 and 1993. Including birth year and province fixed effects should ensure that this does not bias my results. However, the lack of data for these years may lead to relatively imprecise parameter estimates for the first part of my analysis.

(in mm) for each rainfall station, for each month of the 1970-1993 period. I then map each rainfall station to the province (provinsi)<sup>47</sup> where it is located. I average across all stations within a province for each month in the study period. I am left with a data set where the unit of observation is province-year-month rainfall. I create the total rainfall falling in each nine month period, each year, and each three month period by summing across months within a province. Next, I generate running averages for rainfall for each province by averaging over every possible nine-month, year-long, and three-month period. I generate the natural log of each running average (three-month, nine-month and year-long) and the natural log of realized rainfall for every possible three-month, nine-month and year-long period in the data. Thus, the deviation in rainfall for a nine-month period, for example, is calculated by subtracting the natural log of the running average of rainfall for a nine-month period from the natural log of realized rainfall for that nine-month period. This is done separately for each province. A parallel procedure is used to generate deviations for annual and three-month periods.<sup>48</sup> Panel A of Figure 2 plots the distribution of province level rainfall exposure in utero and province level rainfall exposure during the first year of life for my sample. Panel B does the same for rainfall exposure during the third trimester and the first three months after birth. Each graph also plots a normal distribution for comparison. Panel A of Figure 3 plots the distribution of in utero rainfall exposure by age, for the entire sample. Panel B of Figure 3 does the same for province level rainfall exposure during the first year after birth. Both rainfall in utero and rainfall during the first year of life appear close to normally distributed.<sup>49</sup> Rainfall for three month periods appears to be approximately log normally

---

<sup>47</sup>In 1993 there were 27 provinces. Since then eight more provinces have been created. Provinces are the highest tier of local country government subdivisions in Indonesia.

<sup>48</sup>I also generate rainfall deciles and a mean zero, standard deviation one measure of rainfall for each province. The mean zero measure is calculated by subtracting the mean rainfall for a particular province and time period (nine-month or year-long) from the realized rainfall during each possible period and dividing by the standard deviation for that province and time period. I drop any year-province combinations missing more than two months of data and any nine-month-province combinations missing more than one month of data. All of my results are robust to including dummies for the number of months used in calculating the total rainfall for a particular period.

<sup>49</sup>The distribution of rainfall is essentially unchanged when weighting by the number of children born in a province-year-month.



distributed. The mean for in utero rainfall exposure is approximately 1500mm while the mean for rainfall exposure during the year after birth is 2000mm. Figure 3 shows that there is considerable variation in the rainfall children experience, both within a year and across provinces and within a province across years.<sup>50</sup>

My primary explanatory variable is the deviation of the natural log of rainfall in a particular period from the natural log of the running average of rainfall for a period of that length within a province. I use this measure for two reasons. First, this measure is utilized in prior economic research in Indonesia [Maccini and Yang, 2009]. Using a similar measure facilitates comparison of results. Second, as Maccini and Yang [2009] point out when justifying their use of log deviations, this rainfall measure can be interpreted as the percentage deviation from the province level mean during a particular time period. Thus, interpretation of the coefficient on rainfall deviation is relatively straightforward.

## **IFLS**

For person level data I use RAND's Indonesian Family Life Survey. The Indonesian Family Life Survey (IFLS)<sup>51</sup> is a longitudinal household survey containing information on over 66,000 individuals. The first wave of the survey (IFLS1) was fielded in 1993 and is designed to be representative of 83% of the total population of Indonesia.<sup>52</sup> Surveyors in IFLS1 initially contacted 33,081 people across over 7,000 households. Subsequent waves were conducted in 1997 (IFLS2), 1998 (IFLS2+), 2000 (IFLS3), and 2007 (IFLS4). In later waves, attempts were made to contact all of the initial households as well as any new household members. The IFLS contains detailed information about household consumption, education history, work history, and marital history. In waves two through four, respondents between the ages of seven and twenty-four are asked to complete two separate

---

<sup>50</sup>Since all subjects were initially surveyed during the same year (1993), graphing the distribution separately by age is equivalent to graphing the distribution by year.

<sup>51</sup>For a more detailed description of the data see <http://www.rand.org/labor/FLS/IFLS.html>.

<sup>52</sup>Some remote areas are excluded, likely because of the high expected cost of data collection.

tests designed to measure cognitive ability.<sup>53</sup> In all waves health measurements are obtained by trained nurses. The IFLS also contains year, month, and location of birth for all household members. Using the values of birth year, birth month, and household location,<sup>54</sup> I am able to merge the person level characteristics to the rainfall data at the time of birth. My final data set contains the amount of rainfall each child is exposed to during the nine months prior to their birth (by trimester) and in the year after birth.

I focus on two separate samples: children aged three to five in the first wave and children between the ages of six and ten in the first wave. Unfortunately, the IFLS only asks for birth weight and relative size at birth for children under the age of five. Schooling outcomes are not likely to be complete at the time of the fourth survey for these children. Because my main focus is on how parents respond to changes in early childhood conditions, I will primarily be using the younger sample. I do, however, use the older children to check for effects of in utero rainfall on long-term outcomes. For each child between six and ten in the first wave, I construct five separate outcome variables. Years of school completed is constructed by summing the number of years of school the respondent has *completed* by the fourth wave of the survey. Cognitive ability is measured by a Z-score constructed from responses to the cognitive assessment portion of the IFLS4. I calculate the Z-scores separately for the each of two test difficulty levels.<sup>55</sup> Height is the natural log of height in centimeters in wave four. Weight is the natural log of weight in kilograms in wave four.

For children aged three to five during the first wave, I generate three measures intended

---

<sup>53</sup>In wave two, one test measures math ability and a second measures Indonesian language ability. In waves three and four respondents are asked to take a math test and a test to measure general cognitive ability. Additionally, in waves three and four there are two difficulty levels for each test. The easy test is intended for respondents between the ages of seven and thirteen. The hard test is given to respondents between the ages of fourteen and twenty-four.

<sup>54</sup>In practice, the household location (province) at the time of the survey, when children the main sample of children are between three and five, need not be the actual birth province. However, out-of-province migration is extremely uncommon in the sample. For example, less than one percent of one to twenty-three year-old subjects move out of province between 1993 and 1997. Additionally, households with small children are almost certainly less mobile than adolescents and individuals in their early twenties. Thus, while there may be some measurement error introduced by using household location as birth location, it is very unlikely that my results are impacted.

<sup>55</sup>Younger children were given a slightly easier version of the test.

to capture initial health. The IFLS1 contains mother reports of birth weight and relative size at birth. A number of studies in epidemiology and public health suggest that maternal reports of birth weight and other perinatal characteristics are reasonably accurate [Lucia et al., 2006, Adegboye and Heitmann, 2008, Troude et al., 2008]. I generate Z-scores for birth weight and relative size at birth and use these instead of the untransformed values. I also generate a standardized index that combines birth weight and relative size.<sup>56</sup> I do this by summing the Z-scores for birth weight and relative size, subtracting the mean of the sum, and dividing by the standard deviation of the sum.<sup>57</sup> I standardize my measures of relative size and birth weight for ease and consistency in interpretation.

Finally, I construct an indicator variable for whether a child lives in a household that is at high risk for waterborne disease.<sup>58</sup> I construct this measure from two separate survey questions: “What is the main water source for drinking and cooking for this household?” and “Where do the majority of householders go to the toilet?” Children are coded to “High Risk” if their main source of drinking water is well water, rain water, or river/creek water or if the majority of householders go to the bathroom at a shared toilet, a public toilet, in a creek/river/ditch, in a yard/field, or in a sewer. This is a slightly ad-hoc measure of high risk for waterborne illnesses. However, a number of studies in epidemiology and public health suggest that poor quality water supply and poor sanitation conditions are associated with an increased likelihood of contracting six main diseases: Diarrhea, Ascaris, Dracunculosis, Hookworm, Schistosomiasis, and Trachoma [Gadgil, 1998]. The waterborne pathogens associated with these six diseases are most often transmitted through contaminated drinking water. Consumption of river/creek water and use of a non-private bathroom, particularly a

---

<sup>56</sup>Both measures of initial endowment are reported by the mother. Birth weight is in kilograms. Relative size is constructed from the question “In your opinion, compared with other infants, was [...] bigger, smaller or similar in size?”. Neither birth weight nor relative size is a perfect measure of initial endowment. Both have been shown to be somewhat noisy proxies. However, birth weight in particular has been linked to a number of long term outcomes (Almond et al., 2005, Black et al., 2007, Royer, 2009).

<sup>57</sup>This is the procedure outlined in Chetty et al. [2011] and a slight variation of the method used by Kling et al. [2007].

<sup>58</sup>There is also some evidence that rainfall could increase vector-borne diseases (diseases spread by insect vectors such as Malaria, Helminths, and Dengue). See Gubler et al. [2001].

creek/river/ditch, yard/field, or sewer, puts people at an especially high risk of contaminated drinking water. Furthermore, research by the Asian Development Bank suggests that piped water is significantly less likely to be contaminated than the often stagnant water found in wells [ADB, 2006].

I create two variables that capture early measures of parental investment: the number of weeks a child is breastfed and the number of younger siblings in the family. Length of breastfeeding is reported by the mother for all children under five years of age. Breastfeeding requires a considerable investment of time and energy from the mother. In practice, the importance of breastfeeding to the long-term well-being of a child will depend on how mothers use time and energy if they are not breastfeeding, and the nutritional content and price of available substitutes. In developing countries, where the quality of alternative nutrition is almost certainly lower, breastfeeding is likely to be especially important. A large literature in medicine and public health focuses on the benefits of breastfeeding. Though much of this literature relies on observational correlations, a recent randomized trial from Belarus shows that breastfeeding does have a significant, positive impact on cognitive ability [Kramer et al., 2008].

In theory there might be positive cross-sibling spillovers that lead to a positive causal relationship between family size and child outcomes. However, it is generally expected that any potential positive spillover is wiped out by the reduction in parental investment per child resulting from an increase in family size. Existing empirical evidence supports this hypothesis. Black et al. [2010] use a Norwegian data set to test whether an increase in family size has an effect on IQ. Using twins as an instrument for family size, the authors find that an increase in family size significantly reduces child IQ. An increase in birth spacing<sup>59</sup> also appears to have positive effects on educational achievement [Buckles and Munnich, 2012].

I generate a mean zero, standard deviation one, Z-score for the duration of breastfeed-

---

<sup>59</sup>Time between births

ing. I also generate a Z-score for the inverse of the number of siblings. I use the inverse so that, as with breastfeeding duration, higher values can be interpreted as more investment. To do this, I subtract the number of younger siblings from the maximum number found in the data (8). I identify the number of younger siblings using all three subsequent waves of the IFLS. So for children who are two at the baseline, I observe all subsequent childbirth for their mothers between the child's birth and age seventeen. For children who are five at the baseline I observe all younger siblings born until they are age twenty. I subtract the mean and divide by the standard deviation of this inverse number of siblings to get a Z-score. Finally, I generate a combined index by summing the two standardized measures of investment, subtracting the mean of the sum, and dividing by the standard deviation of the sum.

After merging rainfall data with person-level data from the IFLS and requiring that all observations have non-missing controls, I am left with 2,403 children between the ages of six and ten in the first wave (Sample 2) and 1,350 children aged three to five in the first wave (Sample 1). The sample restrictions result in dropping around twenty percent of children aged three to five and thirty percent of children aged six to ten. For Sample 1, none of the basic demographic characteristics are predictive of sample inclusion. However, with such a large proportion of children missing either rainfall data or basic demographic characteristics, I can not rule out the possibility that my results are partially a result of sample selection.

In general, Sample 1 is used for testing how rainfall exposure in utero affects initial child health and subsequent parental investment. Sample 2 is used for illustrating the long-term effects of rainfall exposure. Table 1 presents summary statistics for Sample 1 and Table 2 does the same for Sample 2.

## 5 Empirical Strategy

To test how rainfall exposure in utero affects the initial health endowments of children, subsequent parental investments, and long-term outcomes, I assume the true model for each dependent variable is the following:

$$Y_{c,p,y} = \alpha_E + \beta_E R_{c,p,y} + \gamma_E X_{c,p,y} + \phi_E Z_{c,p,y} + \varepsilon_{c,p,y} \quad (23)$$

The outcome  $Y$ , for child  $c$ , in province  $p$ , in year  $y$  depends on rainfall ( $R$ ) and a vector of controls. More specifically,  $R$  is a vector of rainfall variables capturing the percentage deviation in each trimester of the in utero period and during the year after birth. The vector  $X_{c,p,y}$  includes age fixed effects interacted with child sex, birth month fixed effects, province fixed effects interacted with child sex, a linear trend that is allowed to differ across provinces and child sex, dummies for the mother's age interacted with child sex, a dummy for whether the child is male, and a dummy for whether the child's household resides in an urban setting interacted with whether the child is male. I assume these  $X$ s are exogenous.<sup>60</sup>  $Y_{c,p,y}$  is also likely to be a function of both time-constant and time-varying unobserved family characteristics,  $Z_{c,p,y}$ .<sup>61,62</sup> As these characteristics are unobserved, I am unable to include them in my analysis. Unbiased estimation of  $\beta_E$  will depend on whether after conditioning on the  $X_{c,p,y}$ , the  $Z_{c,p,y}$  are independent of the measures of rainfall,  $R_{c,p,y}$ . In other words, I use OLS to estimate.<sup>63</sup>

$$Y_{c,p,y} = \hat{\alpha}_E + \hat{\beta}_E R_{c,p,y} + \hat{\gamma}_E X_{c,p,y} + \hat{\varepsilon}_{c,p,y} \quad (24)$$

Where the  $\text{plim}_{n \rightarrow \infty} \hat{\beta}_E$  need not be equal to  $\beta_E$ . For  $\hat{\beta}_E$  to be an unbiased estimator of  $\beta_E$ ,

---

<sup>60</sup>One could argue that mother's age is not exogenous. The point estimates are essentially unchanged when omitting mothers age.

<sup>61</sup>Parental ability is one example of what might belong in  $Z_{c,p,y}$ .

<sup>62</sup>The vector of characteristics that matter likely differ for different outcome variables.

<sup>63</sup>I use Huber-White heteroskedasticity-consistent standard errors and cluster at the province level.

it must be true that conditional on the observed  $X'$ s, the unobserved part of the error term which is correlated with outcomes is independent of rainfall deviations. In other words, unbiased estimation requires that the realization of rainfall be uncorrelated with the omitted vector of child characteristics after conditioning on the observed variables  $X_{c,p,y}$ .

This assumption would not be met if unobservably better families were able to predict future rainfall deviations and moved to “wetter” or “drier” provinces. For example, if more educated mothers tend to have children with higher birth weights, have fewer children, and are able to time their pregnancies to periods of higher than average rainfall then the above specification would be biased. However, because my rainfall exposure variables capture deviations in rainfall from the running average of rainfall *within a province*, unobservably better families would have to predict when rainfall will be higher (or lower) than normal for that province. This seems unlikely to be the case. If better parents move to better provinces on average, using deviations in rainfall and including province fixed effects will prevent this from being problematic. In practice, “migration selection” issues are unlikely to be relevant. Movement across Indonesian provinces is very rare in my sample. For example, between 1993 and 1997 less than one percent of one to twenty-three year-olds in the IFLS data move across provinces.

Alternatively, one might worry that unobservably better parents are better at timing their pregnancies to anticipated rainfall or lack thereof. There is some evidence of this in both developing and developed countries [Buckles and Hungerman, 2010, Pitt and Sigle, 2012] though one recent paper [Currie and Schwandt, 2013] suggests that seasonal patterns in the birth outcomes of children in the United States may be the result of the timing of the influenza season. If the months that constitute the rainy season are constant over time, my inclusion of birth month fixed effects and province fixed effects should absorb all of these unobserved factors. However, if the timing of the rainy season changes during the study period and some mothers are better able to time their pregnancies to changes in the timing of the rainy season, the treatment effects I estimate will also include selection bias. There

are three reasons why I believe this is unlikely to affect my results. First, the study period is relatively short. Parents would have to pick up on small changes in the timing of the rainy season over only a seven year period.<sup>64</sup> Second, while there is evidence that parents are able to time births to seasons, the change in the timing of the rainy season, if present, will be very small. Thus, parents would need to be able to time childbirth with a level of precision that is unlikely to be plausible. Finally, including a full set of month of birth by province of birth fixed effects does not lead to statistically different estimates.

A final concern is selective mortality. If the true effect of in utero rainfall exposure is positive, children who experience more rainfall in utero may be more likely to be alive at the time of the IFLS1. The surviving children from provinces exposed to low rainfall would be positively selected. This would bias the coefficients on in utero rainfall downwards.<sup>65</sup> Similarly, if the true effect of exposure to in utero rainfall is negative then children who are exposed to more rainfall in utero would be less likely to survive. Surviving children from provinces exposed to high rainfall would then be positively selected. This would bias the coefficient on in utero rainfall upwards. To test whether there is any selective mortality in my data I collapse all observations by province/year of birth/month of birth. For each cell, I generate a count of the number of children, the rainfall exposure in utero, the rainfall exposure in the year after birth, the rainfall exposure 10 to 18 months before birth, and the rainfall exposure 13 to 24 months after birth. Table 2 displays the coefficients on in utero rainfall exposure from regressions of the count of the number of kids in each cell on rainfall exposure in utero, birth month fixed effects, birth year fixed effects, province fixed effects, and a linear trend for each province. Column 2 includes year of birth rainfall, rainfall 10 to 18 months prior to birth, and rainfall 13 to 24 months after birth. In no column is the coefficient on in utero rainfall statistically significant. Thus, in utero rainfall exposure appears to have no effect on the number of children found in a particular birth year/birth month/province cell. Selective mortality is therefore unlikely to bias my results.

---

<sup>64</sup>I only use children between the ages of three and ten.

<sup>65</sup>This assumes that the outcome variable is constructed so higher values indicate “better” outcomes.



## 6 Results

### OLS Estimates

I begin by exploring the effect of in utero rainfall exposure on initial endowments as measured by birth weight, relative size, and an index combining birth weight and relative size.

I do this by estimating:

$$IE_{c,p,y} = \alpha_E + \beta_E R_{c,p,y} + \gamma_E X_{c,p,y} + \varepsilon_{c,p,y} \quad (25)$$

$IE_{c,p,y}$  are my measures of initial endowments. I estimate (4) on Sample 1, children aged three to five in 1993. Table 3 displays the results. In column 1 the dependent variable is the standardized combined index of early child health. In column 2 the dependent variable is the standardized measure of mother reported birth weight. In column 3 the dependent variable is the standardized index of relative size. All columns include the p-value from a F-test of whether the effect of third trimester rainfall deviation is the same for males and females. Rainfall exposure in utero is associated with significant increases in initial endowments. The point estimates for each trimester are positive, although only the third trimester rainfall deviations are significantly different from zero in each column. The results suggest that a 0.5 log point deviation increase of third trimester rainfall leads to a 0.12 standard deviation increase in the combined initial health index. This effect is significant at the 1% level. Increased third trimester rainfall also has a positive, significant effect on both birth weight and relative size. A 0.5 log point increase in third trimester rainfall is associated with a 0.14 standard deviation increase in birth weight and a 0.10 standard deviation increase in relative size. A 0.14 standard deviation increase in birth weight implies an increase of approximately 0.094 Kg or about 3% of the mean birth weight for the sample. A positive result implies that the income/nutrition effect of third trimester rainfall dominates. Though negative substitution and direct health effects may still be present, they are

outweighed by the increase in available income and nutrition resulting from the increase in rainfall. Figure 4 presents local polynomial smooth plots of standardized birth weight on rainfall deviations in the third trimester for children between the ages of three and five in the first wave of the IFLS. Coinciding with the linear regression output, there is a clear positive relationship between third trimester rainfall deviations and initial health.

Also displayed are the estimated coefficients on year of birth deviation in rainfall. Because mothers are reporting birth weight and relative size at the time of birth, neither measure should be impacted by rainfall during the year after birth. As expected, rainfall deviation in the year after birth is never significantly associated with any of the initial health measures. Similarly, the p-values at the base of each column indicate that I can never reject that the effect of third trimester rainfall is the same for male and female children. This is reassuring given that expectant parents in Indonesia over this time period had little to no access to sex predicting technology. However, the estimates for males, while quantitatively and qualitatively similar to those for females, are noisier despite a similar sample size.

Therefore, rainfall exposure in utero represents a positive shock to initial health endowments for the children in my sample. This is consistent with prior literature which also finds positive effects of in utero rainfall on both short- and long-term outcomes. The finding that the effects of rainfall deviations are only consistently significant in the third trimester is novel.

I now estimate whether long-term outcomes are also affected. To do so, I use Sample 2. Children who are aged three to five in 1993 are likely still in school at the time of the fourth wave. Children ages six to ten in 1993, however, are between twenty-one and twenty-five at the time of IFLS4. It is possible, but unlikely, that they are still in school. Table 4 shows the estimated effect of in utero rainfall deviations by trimester on years of school completed, wave four cognitive ability, and the natural log of wave four height (Panels A, B, and C respectively). Each panel presents the coefficient on the third trimester rainfall deviation, which was previously shown to have a significant and sizable effect on

initial child health. Each column also controls for the first and second trimester rainfall deviations, the rainfall deviation in the year after birth, and the same set of controls as in Table 3. Despite a robust literature in economics linking early conditions to long-term outcomes, Table 4 suggests there is little relationship between exogenous early rainfall shocks and three measures of long-term well-being. None of the results are significant at conventional levels. The point estimate in column 1 suggests that a 0.5 log point increase in third trimester rainfall deviation leads to the completion of roughly .047 fewer years of school for females. To give some context, Duflo [2001] finds that a massive school construction program in Indonesia between 1973 and 1979 led to increases of between 0.25 to 0.40 years of schooling. Using the rapid construction of schools as an instrument for years of schooling, Duflo estimates the return to an extra year of school in Indonesia to be between 3.5% and 10.6%.<sup>66</sup> If I assume the individuals induced to complete fewer years of schooling because of increased rainfall exposure in utero have the same return to education as those induced to complete more years of schooling by school construction<sup>67</sup> this suggests that a 0.5 log point increase in the rainfall deviation in the third trimester leads to a 0.11% to 0.50% reduction in earnings for females. Thus, even ignoring the lack of statistical significance, the point estimates indicate that there is little long-term impact of third trimester rainfall, despite the large and robust effect on early child health. As in Table 3, Table 4 presents the p-values from tests of whether the effect of third trimester rainfall is different for males and females; I can never reject the null that the two effects are the same.

As discussed in the theory section, it is unclear whether this is purely a direct result of rainfall shocks or if there are indirect effects working through other causal pathways. Of particular interest is how parental investment mediates rainfall shocks.<sup>68</sup> Here I provide

---

<sup>66</sup>The wide range is largely a result of different assumptions about the shadow wages for individuals not participating in the wage labor market.

<sup>67</sup>I also need to assume that the rate of return to schooling has not changed since the 1990's. Both of these assumptions are unlikely to be met.

<sup>68</sup>In the treatment effects literature, these two effects are referred to as the average direct effect and the average indirect effect[Huber, 2012]. Specifically, the average direct effect can be written as  $Y(R = r, I(r)) - Y(R = r - 1, I(r))$ . The average indirect effect can be written  $Y(r, I(r)) - Y(r, I(r - 1))$ . Here  $Y(r, i(r))$ , the potential outcome, is a function of realized rainfall and parental investment, which itself is a function of

reduced form evidence on how the potential mediator (parental investment) responds to rainfall shocks.

To test for responses in parental behavior I begin by estimating the following model using the original sample of three to five year-olds in 1993:

$$PI_{c,p,y} = \alpha + \beta R_{c,p,y} + \gamma X_{c,p,y} + u_{c,p,y} \quad (26)$$

Parental investment for child  $c$  in province  $p$  in year  $y$  depends on rainfall ( $R$ ) and a vector of controls. Given the evidence already presented on the positive relationship between third trimester rainfall exposure and initial child health, if parents compensate for early conditions  $\beta$  will be less than zero. If parents reinforce initial shocks  $\beta$  will be greater than zero. I consider three measures of parental investment: the number of younger siblings, the length of breastfeeding, and an index combining the number of younger siblings and the length of breastfeeding. Length of breastfeeding and the combined index are expressed as Z-scores. Table 5 presents the results.

Table 5 presents evidence that parents compensate for initial health shocks. Though not statistically significant for either sex, the point estimates for length of breastfeeding are negative for both males and females. The combined result suggests a 0.5 log point increase in rainfall deviation leads to roughly a 0.014 standard deviation decrease (a 1.5 day decrease) in length of time a child is breastfed.

Both the inverse of the number of younger siblings and the combined index are significantly, negatively impacted for female children but not for male children. That is, an increase in rainfall exposure in utero leads to significantly more younger siblings, but only for female children. Though the p-value at the bottom of columns 8 and 9 suggests the

---

realized rainfall. Obviously only one term in each equation is actually observed. Therefore, identification requires a number of assumptions: that the rainfall measure be as good as randomly assigned and that the potential outcome variable be independent of the investment variable conditional on the realized values of rainfall and any additional covariates included in the model. Future work will explore the feasibility of estimating these effects separately.

difference in the coefficients by child sex is not statistically significant it is nearly three times larger for females than for males. This is also suggestive of compensatory behavior on behalf of parents. The point estimate for females indicates that a 0.5 log point increase in rainfall deviation induces a 0.064 standard deviation increase in the number of younger siblings (decrease in the inverse number of younger siblings) subsequently born to the female child's mother. This corresponds to, on average, 0.08 more younger siblings. The point estimates for males are also negative but much less precise and never statistically significant.

#### IV Estimates

Up to this point, the evidence presented on parental response to initial child health has been indirect. In particular, I first showed that third trimester rainfall exposure has a positive effect on early health and then presented evidence that parents reduce subsequent investment in response to more rainfall, particularly for female children. To provide more direct evidence, I attempt to directly estimate how parents adjust their investment in response to early child health. To do so, I estimate several different specifications with the combined index of parental investment as the dependent variable and the combined index of early health as an explanatory variable. That is, I estimate:

$$PI_{c,p,y} = \alpha + \beta EH_{c,p,y} + \gamma X_{c,p,y} + u_{c,p,y} \quad (27)$$

Where  $EH_{c,p,y}$  is the standardized measure of early health. Note that early health is affected by a number of characteristics that I do not observe in the data. If these characteristics are also correlated with parental investment, OLS estimates of  $\beta$  will be biased. For example,  $X_{c,p,y}$  contains no measure of income, wealth, or parental quality. We expect all three of these unobserved characteristics to be positively correlated with both early health and the measures of parental investment. Thus, the OLS estimate of  $\beta$  is likely to

be positively biased. To deal with this issue, I also estimate instrumental variable specifications, using the rainfall deviation in the third trimester to generate plausibly exogenous variation in early health. Specifically, I estimate limited information maximum likelihood instrumental variables regressions where the first stage is given by:

$$EH_{c,p,y} = \alpha_0 + \theta R_{c,p,y} + \eta X_{c,p,y} + v_{c,p,y} \quad (28)$$

For the IV estimate of  $\beta$  to be consistent, it must be the case that  $E [R'_{c,p,y} EH_{c,p,y}] \neq 0$  and  $E [R'_{c,p,y} u_{c,p,y}] = 0$ .<sup>69</sup> That is, I need that rainfall deviations affect early child health and that conditional on the  $X$ 's, rainfall deviations only affect parental investment through their effect on early health. The latter assumption will be violated if all of the benefits of third trimester rainfall exposure are not captured by my measures of initial health. For example, if mothers receive higher wages and are able to reduce the amount of labor they supply in future periods, allowing more time for breastfeeding, then the instrument would be correlated with the error term in the outcome equation. However, as mentioned earlier, households in developing countries frequently lack the ability to shift income across periods. Of particular relevance are the results of Jayachandran [2006], which indicate that agricultural labor supply is extremely inelastic. Also important is the extensive literature in development suggesting that many households are unable to smooth consumption over time. Therefore, I feel comfortable making the assumption that rainfall deviations are not correlated with the error term in the investment equation.

Because the OLS estimates from the previous section suggest there may be differences across child gender in how parents respond, I would ideally instrument for two endogenous variables: early health and the interaction between early health and a dummy for being male. Unfortunately, the “first-stage” is too weak for male children. Therefore, in all IV estimates, I pool across child sex. Estimating parameters separately by child sex or with

---

<sup>69</sup>At least in a homogeneous effects world. With heterogeneous effects a monotonicity assumption would also have to be satisfied.

an interaction on the endogenous variable and instrument yields point estimates that are statistically indistinguishable across child sex.

Table 6 presents OLS and IV estimates of parental investment responses to early health. Columns 1, 2 and 3 display the estimated coefficients for rainfall deviation in the third trimester from OLS specifications for Males and Females, Females only and Males only, respectively. All columns include the full set of  $X$ s discussed in previous regressions.

As expected, the OLS estimates appear to be positively biased. When pooling across child sex and for males and females separately, the OLS estimates are fairly precise zeros. However, The IV estimate suggests that a one standard deviation increase in early child health<sup>70</sup> causes a 0.466 standard deviation decrease in parental investment. To offer some context, a 0.466 standard deviation decrease in breastfeeding represents a seven week reduction in breastfeeding length. A 0.466 standard deviation decrease in the number of younger siblings represents a decrease of 0.60 younger siblings, on average. These estimates indicate there are economically important reductions (increases) in parental investment in children in response to positive (negative), exogenous variation in early health.<sup>71</sup>

### Testing The Model Prediction

The model in Section 3 offered one testable prediction:

1. The likelihood that  $\frac{\partial i^*(\varepsilon)}{\partial \varepsilon} > 0$  is increasing in the price of investment.

As mentioned above, an increase in rainfall should increase the return to supplying labor or capital to agricultural production. Thus, an increase in rainfall should increase the opportunity cost of investment in children. One period during which this may be particularly important is shortly after birth when the human capital of the child is still highly malleable. However, as with in utero rainfall exposure, we might also expect there to be a direct, positive income effect of rainfall after birth.<sup>72</sup> That said, after conditioning on

---

<sup>70</sup>An increase that is induced by the rainfall deviation in the third trimester of pregnancy.

<sup>71</sup>The excluded instruments in the IV estimate presented in column 4 have a Kleibergen-Paap F-statistic of 10.452, well above the critical values for a LIML test of 15% size outlined in Stock and Yogo [2002].

<sup>72</sup>See Maccini and Yang [2009] and Millett and Shah [2012a].

rainfall exposure during the year after birth and early child health independently, it is difficult to see why their interaction would impact parental investment, except by increasing the opportunity cost of investment. A positive coefficient on this interaction term would then confirm the prediction of the model.

To test this prediction I generate a dummy variable that is equal to one if the rainfall deviation a child experiences during the twelve months after birth is above the median for the sample. I then interact this dummy with the rainfall deviation the child experienced during the third trimester of pregnancy. Ideally, this test would be implemented using the IV specification used in Section 6.0.2 with initial child health interacted with post-birth rainfall exposure for both females and males. Unfortunately, I do not have sufficient first-stage power to confidently instrument for two, let alone four different endogenous variables. Therefore, I estimate this specification, separately for females and males and pooling across child sex using OLS. Importantly, I condition on both the non-interacted dummy variable for above median year after birth rainfall and a continuous measure of the rainfall deviation in the year after birth.<sup>73</sup> Together, these should capture any income effect resulting from post-birth rainfall. In results not shown, I pool across child sex and interact the endogenous variable (early health) and instrument (rainfall deviation in the third trimester) with the indicator for whether the rainfall deviation in the year after birth is above the median and estimate liml IV models. Though noisy, IV estimates are signed the same as the OLS estimates discussed below.

Table 7 presents results. For both females and males the interaction term is positive and significant. This suggests that parents are more likely to reinforce early conditions when the opportunity cost of investment in children is higher as the model predicts. The point estimates indicate that the investment response to the interaction between post-birth rainfall and third trimester rainfall is slightly larger for females than for males. However, I fail to reject that the effect is the same across sex and the estimates are positive and significant

---

<sup>73</sup>I include the same set of controls used in the estimates from Section 6.0.1.



for both males and females. For females, a log point increase in rainfall deviation during the third trimester leads to a 0.216 standard deviation decrease in parental investment if the child experienced below median rainfall in the year after birth but only a 0.047 standard deviation decrease if rainfall was above the median in the year after birth. The corresponding effects for males are a 0.119 standard deviation decrease for below median rainfall children and a .027 standard deviation *increase* for above median rainfall children.

To help illustrate the model prediction tested above consider two female children who come from identical households and are exposed to the same negative, third trimester rainfall shock, but only one of is exposed to above median year after birth rainfall. The model and data suggest that both children will be born smaller and lighter as a result of the rainfall shock. However, while both sets of parents will compensate for this early deficit through increased investment (an increase in duration of breastfeeding and a decrease in the number of younger siblings), the child who experiences less rainfall during the year after birth will benefit from significantly more compensating investment. Under the assumption that the controls eliminate all potentially confounding factors this is evidence that compensating behavior is decreasing in the price of investment and reinforcing behavior is increasing in the price of investment.

## **7 Conclusion**

This paper presents evidence that an increase in exposure to in utero rainfall has significant short- and medium-run effects on individual well-being in Indonesia. Children exposed to positive rainfall deviations in the third trimester of pregnancy are born significantly heavier and significantly larger than their non-exposed peers. This suggests that increases in rainfall may increase initial endowments through a positive income effect. Despite these early advantages, long-term outcomes do not appear to be positively impacted by in utero rainfall exposure. Differences in parental investment are one potential mechanism for explaining

these results. IV estimates suggest that girls who were exposed to positive rainfall shocks in utero have a greater number of younger siblings and were breastfed for less time. On the other hand, the response in parental investment for boys, while negatively signed, is imprecise, not significantly different from zero and qualitatively smaller than for girls. Together these results suggest that parents in Indonesia compensate initial health conditions and that this pattern is more pronounced for girls than for boys. More generally, the results support the idea that changes in parental investment can remediate early health deficits and nullify early health advantages.

## 8 Tables

**Table 1**

Panel A: Summary Statistics Children Ages Three to Five in 1993						
General Characteristics:	mean	sd	min	p50	max	count
Age in 1993	3.986	0.809	3.000	4.000	5.000	1350
Male	0.509	0.500	0.000	1.000	1.000	1350
Lives in and Urban Setting in 1993	0.449	0.498	0.000	0.000	1.000	1350
Age of Mother in 1993	31.044	6.657	16.000	30.000	56.000	1350
Month of Birth	6.582	3.306	1.000	7.000	12.000	1350
Rainfall Variables:						
Rain YOB (mm)	1681.692	615.011	204.500	1661.000	3520.000	1350
Rain In UT (mm)	1421.256	443.632	311.000	1366.100	2956.917	1350
Rain Third Trimester (mm)	459.312	291.910	1.000	453.100	2085.833	1350
Rain Second Trimester (mm)	488.333	295.265	1.000	474.100	2085.833	1350
Rain First Trimester (mm)	473.611	286.121	2.000	470.907	2085.833	1350
Rainfall Deviation YOB	-0.209	0.442	-2.081	-0.125	0.690	1350
Rainfall Deviation In UT	-0.061	0.307	-1.313	-0.067	0.866	1350
Rainfall Deviation Third Trimester	-0.373	0.920	-5.596	-0.185	1.267	1350
Rainfall Deviation Second Trimester	-0.268	0.836	-5.596	-0.099	1.267	1350
Rainfall Deviation First Trimester	-0.309	0.852	-4.902	-0.099	1.310	1350
Early Health and Parental Investment:						
Birth Weight (Kg)	3.167	0.660	0.300	3.030	7.000	508
Inverse Relative Size	2.171	0.725	0.000	2.000	4.000	883
Length of Time Breastfed (Weeks)	15.780	15.520	0.000	13.000	52.000	1157
Number of Younger Siblings	1.395	1.284	0.000	1.000	6.000	1350

Note: Table 1 Panel A presents summary statistics for children matched to rainfall measures and ages three to five in 1993. Observations decrease for birth weight and relative size because some mothers elected not to respond to these questions. Rainfall deviations are calculated as  $\ln(\text{rain in period } x) - \ln(\text{RA period length } x)$  where RA period length  $x$  represents the running average over the entire sample for a period of length  $x$ . Rainfall deviations have negative means because rainfall for three month periods is distributed approximately log normally. This should not be problematic for the analysis since rainfall deviations are being compared only across children in this sample.

**Table 1**

Panel B: Summary Statistics Children Ages Six to Ten in 1993						
General Characteristics:	mean	sd	min	p50	max	count
Age in 1993	8.023	1.474	6.000	8.000	10.000	1980
Male	0.498	0.500	0.000	0.000	1.000	1980
Lives in and Urban Setting in 1993	0.480	0.500	0.000	0.000	1.000	1980
Age of Mother in 1993	34.518	7.033	20.000	34.000	70.000	1980
Month of Birth	6.578	3.328	1.000	7.000	12.000	1980
Rainfall Variables:						
Rain YOB (mm)	1932.813	453.941	239.000	1926.667	3916.000	1980
Rain In UT (mm)	1454.785	430.652	327.500	1443.667	2868.000	1980
Rain Third Trimester (mm)	481.737	279.005	0.500	474.700	1716.000	1980
Rain Second Trimester (mm)	496.684	287.477	0.500	490.286	2085.833	1980
Rain First Trimester (mm)	476.364	292.020	0.500	472.714	1741.000	1980
Rainfall Deviation YOB	-0.020	0.198	-1.314	-0.038	0.683	1980
Rainfall Deviation In UT	-0.057	0.322	-1.344	-0.043	0.866	1980
Rainfall Deviation Third Trimester	-0.325	0.907	-7.069	-0.120	1.334	1980
Rainfall Deviation Second Trimester	-0.324	1.022	-7.069	-0.062	1.334	1980
Rainfall Deviation First Trimester	-0.409	1.104	-7.069	-0.131	1.334	1980
Long-Term Outcomes:						
Years of School Completed Wave 4	11.046	3.800	0.000	13.000	20.000	1966
Z Score Cognitive Test Wave 4	-0.008	1.026	-3.182	0.224	1.665	1948
Ln Weight(Kg) Wave 4	3.948	0.186	1.569	3.942	4.653	1836
Ln Height(Cm) Wave 4	5.051	0.117	2.815	5.059	5.209	1838

Note: Table 1 Panel B presents summary statistics for children matched to rainfall measures and ages six to ten in 1993. Rainfall deviations are calculated as  $\ln(\text{rain in period } x) - \ln(\text{RA period length } x)$  where RA period length  $x$  represents the running average over the entire sample for a period of length  $x$ . Rainfall deviations have negative means because rainfall for three month periods is distributed approximately log normally. This should not be problematic for the analysis since rainfall deviations are being compared only across children in this sample.

**Table 2**  
Rainfall and Selective Mortality

VARIABLES	Number of Children Born	Number of Children Born
	(1)	(2)
Rainfall During the 9 Months Prior to Birth	0.00006 (0.00015)	0.00002 (0.00015)
Constant	-1.55285 (1.24324)	-2.10794 (1.29280)
Observations	2,380	2,380
R-squared	0.70829	0.70792

Notes: Both columns present coefficients from OLS regressions with the count of children born in a birth year-birth month-province appearing in IFLS 1 as the dependent for a nine month period within a province. All columns also include birth year fixed effects, birth month fixed effects, and linear trends for each province. Column 2 includes afor similarly defined rainfall deviations ten to eighteen months prior to birth and thirteen to twenty-four months after birth. Robust standard errors in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.01

**Table 3**

VARIABLES	Rainfall by Trimester and Initial Health		Relative Size (3)
	Combined Initial Health Index (1)	Birth Weight (2)	
Rainfall Deviation Third Trimester	0.235*** (0.070)	0.286*** (0.064)	0.208** (0.082)
Rainfall Deviation Second Trimester	0.052 (0.063)	0.165* (0.094)	0.038 (0.057)
Rainfall Deviation First Trimester	0.019 (0.075)	0.004 (0.099)	0.013 (0.085)
Rainfall Deviation Year After Birth	0.023 (0.109)	-0.058 (0.075)	0.045 (0.115)
Constant	-0.348 (0.374)	-0.797 (0.318)	0.0180 (0.400)
Observations	886	508	883
Rsquared	0.213	0.391	0.172
Rainfall is Equal For Males and Females	0.730	0.577	0.907

Note: Table 3 presents estimates of the average partial effect from OLS regressions of measures of initial child health on percentage deviation in rainfall during each of the three trimesters of pregnancy, percentage deviation during the year after the child's birth, mother's age, mother's age by child sex, birth year fixed effects, birth year by child sex fixed effects, birth month fixed effects, trends for each province that are allowed to differ by child sex, a dummy for whether the child is male, a dummy for whether the child lived in an urban setting in 1993, the interaction between the male and urban dummies, and a dummy indicating whether the child was missing either birth weight or relative size. In addition, columns also display the p-value resulting from a test of whether the average partial effect of third trimester rainfall is equal across child sex when rainfall deviation by child sex interactions are included. The sample is limited to children between the ages of three and five in 1993. Each regression is weighted using the IFLS longitudinal (Wave 4) weights. Standard errors are clustered at the province level. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.

**Table 4**

VARIABLES	Third Trimester Rainfall and Long-Term Outcomes		
	Wave 4	Wave 4	Wave 4
	Years of Schooling Completed (1)	Z Score Cognitive Ability (2)	Ln Height (CM) (3)
Rainfall Deviation Third Trimester	-0.094 (0.127)	0.041 (0.038)	0.000 (0.003)
Rainfall Deviation Year After Birth	0.647 (0.574)	0.123 (0.168)	-0.039* (0.023)
Observations	1,966	1,948	1,838
Constant	9.582 (0.466)	0.248 (0.206)	5.004 (0.0240)
Rsquared	0.231	0.153	0.152
P-Value: Effect of Third Trimester Rainfall is Equal For Males and	0.214	0.258	0.714

Note: Table 4 presents estimates of the average partial effect of the third trimester rainfall deviation and the rainfall deviation in the year after birth from OLS regressions of measures of Wave 4 outcomes on percentage deviation in rainfall during each of the three trimesters of pregnancy, percentage deviation during the year after the child's birth mother's age, mother's age by child sex, birth year fixed effects, birth year by child sex fixed effects, birth month fixed effects, trends for each province that are allowed to differ by child sex, a dummy for whether the child is male, a dummy for whether the child lived in an urban setting in 1993, and the interaction between the male and urban dummies. In addition, columns also display the p-value resulting from a test of whether the average partial effect of third trimester rainfall is equal across child sex when rainfall deviation by child sex interactions are included. The sample is limited to children between the ages of six and ten in 1993. Each regression is weighted using the IPLS longitudinal (Wave 4) weights. Standard errors are clustered at the province level. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.

**Table 5**

VARIABLES	Third Trimester Rainfall and Parental Investment								
	Combined Parental Investment			Length of Breastfeeding			Inverse Number of Younger		
	Males and Females (1)	Females (2)	Males (3)	Males and Females (4)	Females (5)	Males (6)	Males and Females (7)	Females (8)	Males (9)
Rainfall Deviation Third Trimester	-0.090* (0.052)	-0.121** (0.055)	-0.060 (0.094)	-0.028 (0.045)	-0.035 (0.061)	-0.027 (0.080)	-0.083 (0.057)	-0.127*** (0.041)	-0.036 (0.100)
Rainfall Deviation Year After Birth	-0.199*** (0.066)	-0.096 (0.153)	-0.301* (0.158)	-0.406*** (0.082)	-0.369*** (0.113)	-0.434** (0.178)	0.072 (0.065)	0.219* (0.130)	-0.090 (0.119)
Observations	1,350	1,350	1,350	1,157	1,157	1,157	1,350	1,350	1,350
Constant	0.110 (0.284)	0.0380 (0.302)		-0.245 (0.446)	-0.281 (0.462)		0.351 (0.239)	0.280 (0.226)	
Rsquared	0.230	0.235		0.197	0.200		0.290	0.294	
P-Value: Effect of Third Trimester Rainfall is Equal For Males and Females			0.565		0.940				0.336

Note: Table 5 presents estimates of the average partial effect from OLS regressions of measures of parental investment on percentage deviation in rainfall during each of the three trimesters of pregnancy, percentage deviation during the year after the child's birth, mother's age, mother's age by child sex, birth year fixed effects, birth year by child sex fixed effects, birth month fixed effects, trends for each province that are allowed to differ by child sex, a dummy for whether the child is male, a dummy for whether the child lived in an urban setting in 1993, and the interaction between the male and urban dummies. In addition, columns 2, 3, 5, 6, 8, and 9 include interactions between each of the rainfall deviations and child sex and the p-value resulting from a test of whether the average partial effect of third trimester rainfall is equal across child sex when rainfall deviation by child sex interactions are included. The sample is limited to children between the ages of three and five in 1993. Each regression is weighted using the IPLS longitudinal (Wave 4) weights. Standard errors are clustered at the province level. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.



**Table 6**  
Parental Investment and Initial Child Health

VARIABLES	OLS Combined Index of Parental Investment		IV Combined Parental Investment Index	
	Males and Females (1)	Females (2)	Males (3)	No Male Interaction (4)
Early Health	-0.032 (0.026)	-0.030 (0.063)	-0.039 (0.060)	-0.466** (0.183)
Early Health*Male				
Observations	886		886	886
R-squared	0.266		0.283	-0.163
P-Value: Effect of Third Trimester Rainfall is Equal For Males and			0.924	

Note: Table 6 presents estimates from OLS and limit IV regressions of measures of the parental investment index on the standardized index of initial child health. All specifications control for the rainfall deviation in the first and second trimesters of pregnancy, percentage deviation during the year after the child's birth, mother's age, mother's age by child sex, birth year fixed effects, birth year by child sex fixed effects, birth month fixed effects, trends for each province that are allowed to differ by child sex, a dummy for whether the child is male, a dummy for whether the child lived in an urban setting in 1993, and the interaction between the male and urban dummies. In addition, column 2 includes an interaction between early child health and child sex. Column 3 displays the p-value from a test of equality of the average partial effects of early childhood health across child sex. IV estimates use the rainfall deviation in the third trimester (column 3) as an instrument. The sample is limited to children between the ages of three and five in 1993. Each regression is weighted using the IFLS longitudinal (Wave 4) weights. Standard errors are clustered at the province level. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.

**Table 7**  
Compensating Behavior and The Opportunity Cost of Investment

VARIABLES	OLS		
	Males and Females (1)	Females (2)	Males (3)
Rainfall Deviation Third Trimester	-0.144** (0.058)	-0.216** (0.087)	-0.119 (0.103)
Rainfall Deviation Third Trimester* Above Median Year After Birth	0.122**	0.169* (0.095)	0.146** (0.073)
Rainfall Deviation	-0.059		
Observations	1,350	1,350	
R-squared	0.230		0.236
P-Value: Effect of Third Trimester Rainfall is Equal For Males and Females			0.495
P-Value: Interaction between Third Trimester Rainfall and Above Median Year After Birth Rainfall Deviation is Equal For Males and Females			0.884

Note: Table 7 presents estimates from OLS regressions of measures of the parental investment index on the rainfall deviation during the third trimester of pregnancy (Columns 1, 2 and 3) and their interaction with an indicator variable for whether rainfall in the year after birth was above the median rainfall for a year in the province of birth. All specifications control for the rainfall deviation in the first and second trimesters of pregnancy, percentage deviation during the year after the child's birth, mother's age, mother's age by child sex, birth year fixed effects, birth year by child sex fixed effects, birth month fixed effects, trends for each province that are allowed to differ by child sex, a dummy for whether the child is male, a dummy for whether the child lived in an urban setting in 1993, an indicator for whether the rainfall deviation in the year after birth was above the median for that province and the interaction between the male and urban dummies. The sample is limited to children between the ages of three and five in 1993. Each regression is weighted using the IFLS longitudinal (Wave 4) weights. Standard errors are clustered at the province level. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1.

## 9 Figures

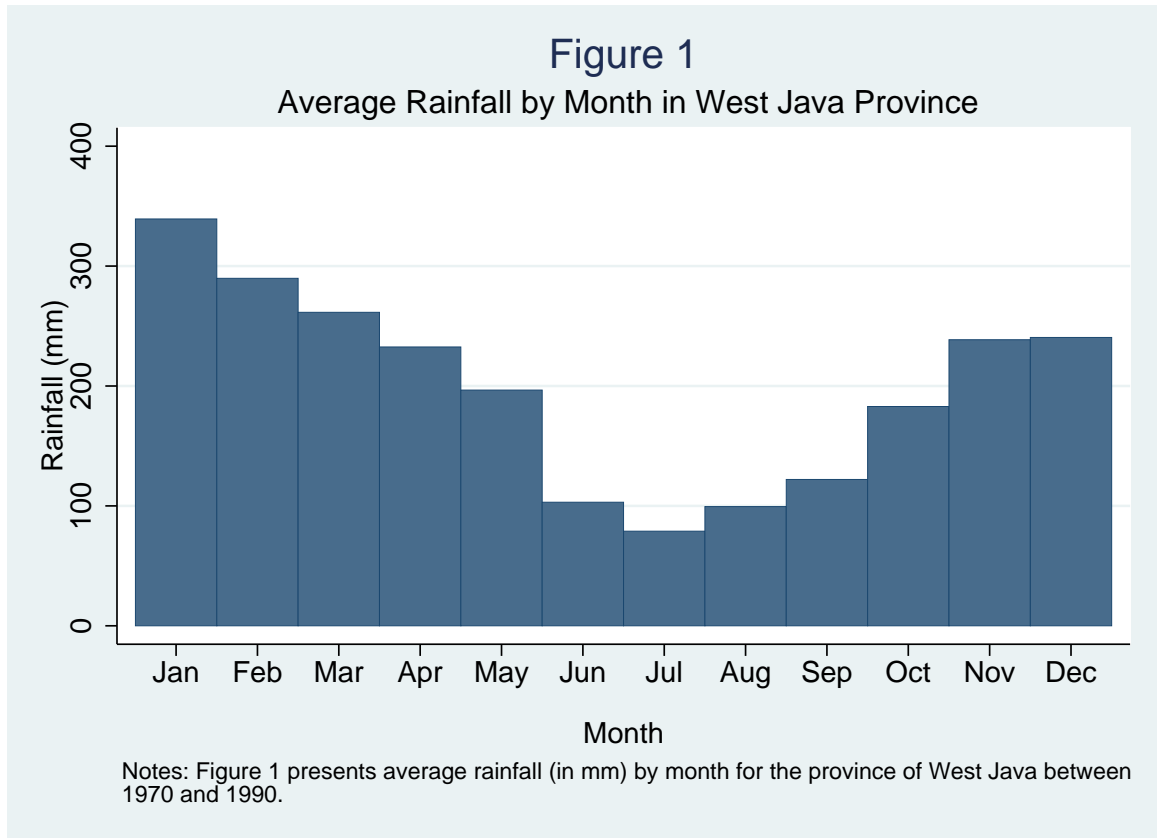
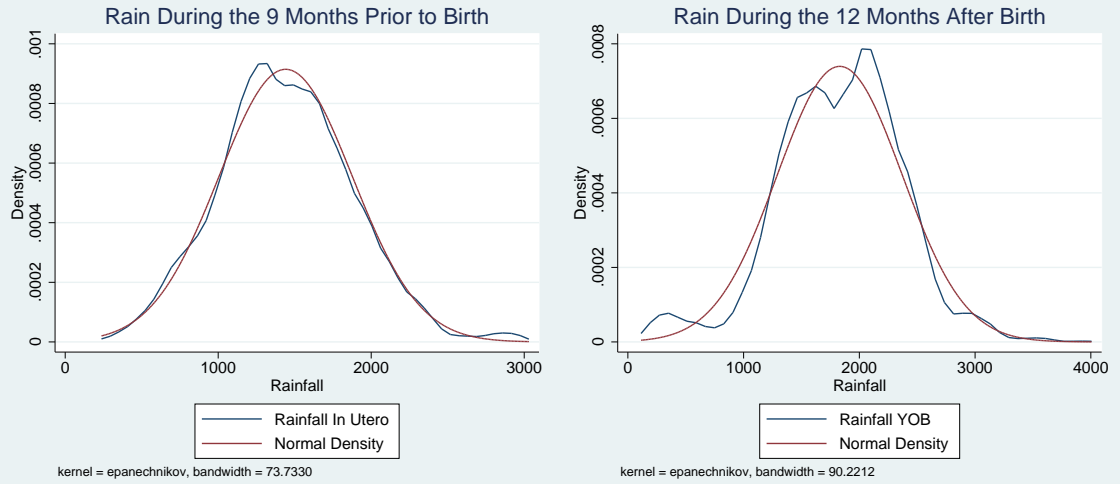
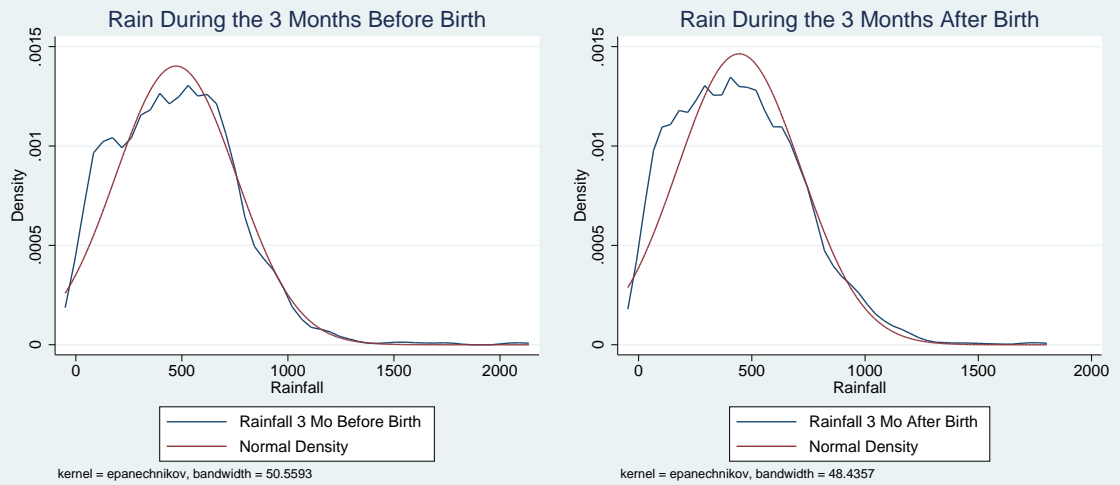


Figure 2: Panel A



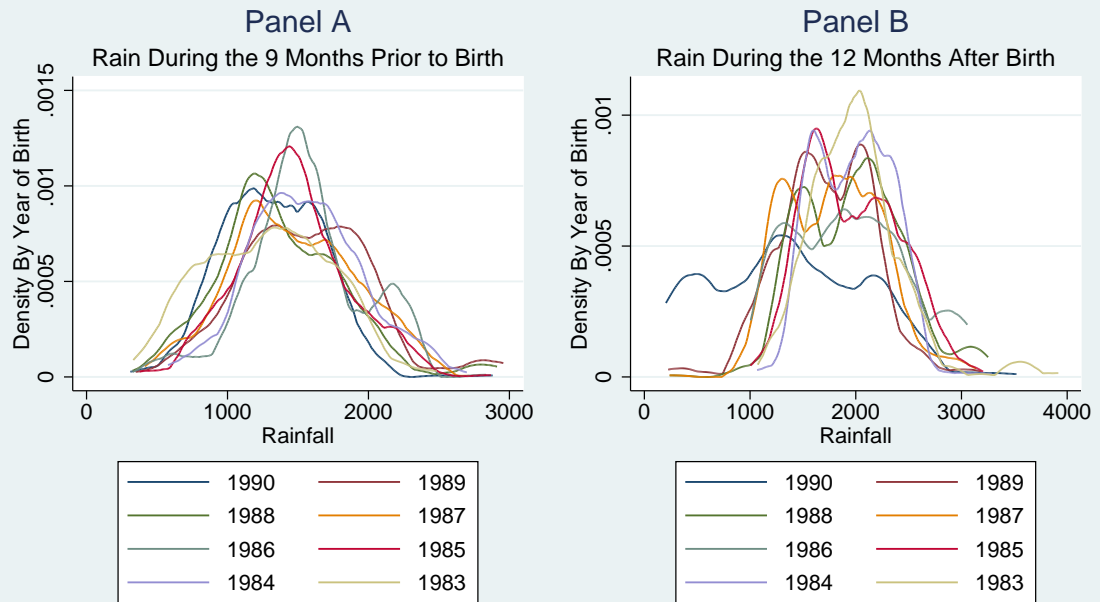
Notes: Figure 2 panel A presents a kernel density plot of rainfall (in mm) during the nine months prior to birth in the province of birth for children age 3 to 5 in 1993.

Figure 2: Panel B



Note: Figure 2 panel B presents a kernel density plot of rainfall (in mm) during the year after birth in the province of birth for children age 3 to 5 in 1993.

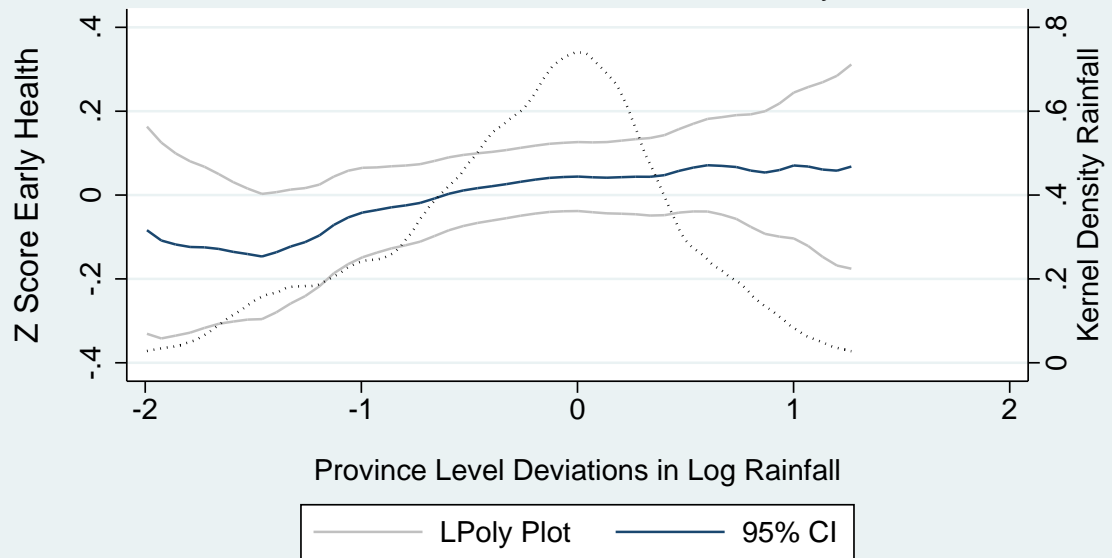
Figure 3



Note: Panel A presents kernel density plots of rainfall (in mm) during the nine months prior to birth in the province of birth by year of birth for children age 3 to 10 in 1993. Panel B presents kernel density plots of rainfall (in mm) during the year after birth in the province of birth by year of birth for children age 3 to 10 in 1993.

Figure 4

Third Trimester Rainfall Deviation and Early Health



Note: Figure 4 presents results from local polynomial regression of the standardized index of early child health on the rainfall deviation experienced by the child during the three months prior to birth. The estimates employ an Epanechnikov kernel and the rule-of-thumb plugin bandwidth. The 95% confidence interval is represented by the gray lines. The density of the rainfall deviation variable is presented using a dotted line.

## 10 Appendix A: Model

$$\max_{i_s, c} u(q_s, c) \text{ s.t. } y \geq pi_s - c, q_s = f(i_s, \varepsilon_s) \quad (29)$$

Where  $q_s, s \in \{f, m\}$ , represents the quality for a child of sex  $s$ , and  $c$  is a numeraire good.  $q_s$  is determined by  $q_s = f(i_s, \varepsilon_s)$  where  $\varepsilon_s$  captures exogenous, initial conditions and  $i_s$  is a measure of parental investment. Assuming Inada conditions on the utility and production functions<sup>74</sup>, first order conditions imply that at the solution:

$i_s$  :

$$u_q(f(i_s^*, \varepsilon_s), c^*) f_{i_s}(i_s^*, \varepsilon_s) - \lambda p = 0 \quad (30)$$

$c$  :

$$u_c(f(i_s^*, \varepsilon_s), c^*) - \lambda = 0 \quad (31)$$

$\lambda$  :

$$y = pi_s^* - c^* \quad (32)$$

Combining the first two conditions yields:

$$pu_c(f(i_s^*, \varepsilon_s), c^*) = u_q(f(i_s^*, \varepsilon_s), c^*) f_{i_s}(i_s^*, \varepsilon_s) \quad (33)$$

Which offers the usual intuition. Parents invest in child quality until the increase in utility from an extra unit of investment is equated to the price scaled increase in utility from an extra unit of consumption of the numeraire good. In other words, parents consume until the marginal rate of substitution between the two goods equals the price ratio. For my purposes, I am interested in how parents adjust their optimal investment in response to changes in initial child quality. That is, I care about signing  $\frac{\partial i_s^*(\varepsilon)}{\partial \varepsilon}$ . To do so, I rewrite the two remaining first order conditions as functions of the choice variables and the initial child condition and apply Cramer's rule to the resulting system. See Appendix A for details.

---

<sup>74</sup> $\lim_{x \rightarrow 0} u_c(q_s, x) = +\infty, \lim_{z \rightarrow 0} u_q(z, c) = +\infty, u_q > 0, u_c > 0, u_{qq} < 0, u_{cc} < 0, f(0, \varepsilon) = f(i, 0) = 0, f_{ii} < 0, f_{\varepsilon\varepsilon} < 0, f_i > 0, f_\varepsilon > 0$

After:

$$F_1(c^*, i_s^*; \varepsilon_s) = pu_c(f(i_s^*(\varepsilon_s), \varepsilon_s), c^*(\varepsilon_s)) - u_q(f(i_s^*(\varepsilon_s), \varepsilon_s), c^*(\varepsilon_s)) f_{i_s}(i_s^*(\varepsilon_s), \varepsilon_s) = 0 \quad (34)$$

$$F_2(c^*, i_s^*; \varepsilon_s) = y - pi_s^*(\varepsilon) - c^*(\varepsilon) = 0 \quad (35)$$

Applying the chain rule and writing the result in matrix form yields:

$$\begin{bmatrix} \frac{\partial F_1}{\partial c^*(\varepsilon)} & \frac{\partial F_1}{\partial i_s^*(\varepsilon)} \\ \frac{\partial F_2}{\partial c^*(\varepsilon)} & \frac{\partial F_2}{\partial i_s^*(\varepsilon)} \end{bmatrix} \begin{bmatrix} \frac{\partial c^*(\varepsilon)}{\partial \varepsilon} \\ \frac{\partial i_s^*(\varepsilon)}{\partial \varepsilon} \end{bmatrix} = - \begin{bmatrix} \frac{\partial F_1}{\partial \varepsilon} \\ \frac{\partial F_2}{\partial \varepsilon} \end{bmatrix} \quad (36)$$

Applying Cramer's rule to the resulting system provides the following expression for  $\frac{\partial i_s^*(\varepsilon)}{\partial \varepsilon}$ :

$$\frac{\partial i_s^*(\varepsilon)}{\partial \varepsilon} = - \frac{\det \begin{bmatrix} \frac{\partial F_1}{\partial c^*(\varepsilon)} & \frac{\partial F_1}{\partial \varepsilon} \\ \frac{\partial F_2}{\partial c^*(\varepsilon)} & \frac{\partial F_2}{\partial \varepsilon} \end{bmatrix}}{\det \begin{bmatrix} \frac{\partial F_1}{\partial c^*(\varepsilon)} & \frac{\partial F_1}{\partial i_s^*(\varepsilon)} \\ \frac{\partial F_2}{\partial c^*(\varepsilon)} & \frac{\partial F_2}{\partial i_s^*(\varepsilon)} \end{bmatrix}} \quad (37)$$

Expanding the determinant expressions:

$$\frac{\partial i_s^*(\varepsilon)}{\partial \varepsilon} = - \frac{\frac{\partial F_1}{\partial c^*(\varepsilon)} \frac{\partial F_2}{\partial \varepsilon} - \frac{\partial F_2}{\partial c^*(\varepsilon)} \frac{\partial F_1}{\partial \varepsilon}}{\frac{\partial F_1}{\partial c^*(\varepsilon)} \frac{\partial F_2}{\partial i_s^*(\varepsilon)} - \frac{\partial F_2}{\partial c^*(\varepsilon)} \frac{\partial F_1}{\partial i_s^*(\varepsilon)}} \quad (38)$$

Each of the terms in this expression can be calculated using relatively simple algebra<sup>75</sup>.

$$\frac{\partial F_1}{\partial c^*(\varepsilon)} = pu_{cc}(f(\cdot), c^*(\cdot)) - u_{qc}(f(\cdot), c^*(\cdot)) f_{i_s}(\cdot) \quad (39)$$

$$\frac{\partial F_1}{\partial i_s^*(\varepsilon)} = pu_{cq}(f(\cdot), c^*(\cdot)) f_{i_s^*}(\cdot) - u_{qq}(f(\cdot), c^*(\cdot)) f_{i_s}(\cdot)^2 - u_q(f(\cdot), c^*(\cdot)) f_{i_s i_s}(\cdot) \quad (40)$$

$$\frac{\partial F_1}{\partial \varepsilon} = pu_{cq}(f(\cdot), c^*(\cdot)) f_{\varepsilon_s}(\cdot) - u_{qq}(\cdot, c^*(\cdot)) f_{\varepsilon_s}(\cdot) f_{i_s}(\cdot) - u_q(f(\cdot), c^*(\cdot)) f_{i_s \varepsilon_s}(\cdot) \quad (41)$$

<sup>75</sup>I suppress the interior of the different parentheses to simplify the display



$$\frac{\partial F_2}{\partial c^* (\varepsilon)} = -1 \quad (42)$$

$$\frac{\partial F_2}{\partial i_s^* (\varepsilon)} = -p \quad (43)$$

$$\frac{\partial F_2}{\partial \varepsilon} = 0 \quad (44)$$

Substituting the expressions into (1):

$$= \frac{0 - (-1) [pu_{cq} (f(\cdot), c^*(\cdot)) f_{\varepsilon_s}(\cdot) - u_{qq} (f(\cdot), c^*(\cdot)) f_{\varepsilon_s}(\cdot) f_{i_s}(\cdot) - u_q (f(\cdot), c^*(\cdot)) f_{i_s \varepsilon_s}(\cdot)]}{[pu_{cc} (f(\cdot), c^*(\cdot)) - u_{qc} (f(\cdot), c^*(\cdot)) f_{i_s}(\cdot)] (-p) - (-1) [pu_{cq} (f(\cdot), c^*(\cdot)) f_{i_s}^* (\cdot) - u_{qq} (f(\cdot), c^*(\cdot)) f_{i_s}(\cdot)^2 - u_q (f(\cdot), c^*(\cdot)) f_{i_s i_s}(\cdot)]} \quad (45)$$

$$= \frac{[pu_{cq} (f(\cdot), c^*(\cdot)) f_{\varepsilon_s}(\cdot) + u_{qq} (f(\cdot), c^*(\cdot)) f_{\varepsilon_s}(\cdot) f_{i_s}(\cdot) + u_q (f(\cdot), c^*(\cdot)) f_{i_s \varepsilon_s}(\cdot)]}{p [pu_{cc} (f(\cdot), c^*(\cdot)) - u_{qc} (f(\cdot), c^*(\cdot)) f_{i_s}(\cdot)] - [pu_{cq} (f(\cdot), c^*(\cdot)) f_{i_s}^* (\cdot) + u_{qq} (f(\cdot), c^*(\cdot)) f_{i_s}(\cdot)^2 + u_q (f(\cdot), c^*(\cdot)) f_{i_s i_s}(\cdot)]} \quad (46)$$

This expression captures the prediction from this simple model for how parents investment behavior responds to early conditions.

## **Part III**

# **Childhood Housing and Adult Earnings: A Between-Siblings Analysis of Housing Vouchers and Public Housing**

## **1 Introduction**

In the year 2000, over 2.7 million children under the age of eighteen lived in voucher-supported or public housing, the two most popular subsidized housing programs run by the Department of Housing and Urban Development (HUD). Although large-scale assisted housing programs have been in place for some time, research on the long-term effects for resident children is scarce, and hampered by methodological and data limitations.

This paper estimates the causal effect of participation in voucher-supported and public housing as a teenager on employment and earnings in early adulthood. To do so, we develop a novel data set that combines information on housing assistance, earnings, household structure, and neighborhood and demographic characteristics. By linking these different data sources together at the person level, we are able to track millions of children as they progress through voucher-supported, public, and unassisted housing as children, and into the labor market as adults.

There are a number of channels through which childhood participation in subsidized housing might impact adult outcomes. Both voucher and public housing provide a positive income effect for households. By expanding the budget set faced by participating households, these programs may enable parents to devote more time and financial resources to develop the human capital of children residing the household. This increase in human cap-

ital should be reflected in higher labor market earnings, suggesting that assisted housing residence in childhood would positively impact adult labor market outcomes.

However, other pathways would yield a negative relationship between subsidized housing participation in childhood and adult labor market performance. Oreopoulos [2003] raises the possibility that subsidized housing participation might impact outcomes through peer or neighborhood effects. If, as argued by Oreopoulos [2003] and Newman [1972], available subsidized housing units are located in worse neighborhoods—i.e. neighborhoods with higher crime rates and lower quality schools—than participants' counterfactual housing options, then public and voucher-assisted housing could have negative neighborhood and peer effects and therefore decrease adult earnings. Ex-ante, the sign of any neighborhood or peer effects, as well as the overall impact of subsidized housing participation, is unclear. Our results identify the net long-term effect on adult earnings of childhood participation in subsidized housing.

Implicitly assumed in the previous paragraphs was the idea that the impact of housing vouchers and public housing participation during childhood is the same. This need not be the case. In fact, the thought that the two programs might have different effects is one element underlying the shift in subsidized housing policy in the U.S. to provide housing choice through vouchers. The argument is that in the absence of discrimination on the part of potential landlords, voucher housing should offer households increased neighborhood choice. As such, any adverse consequences of public housing projects could be avoided while the positive income effect for households would still be present. The debate about housing vouchers vs. public housing has been the subject of previous research, but most of it has focused on the difference in short-term effects. For example, recent evidence indicates that female youth moving to lower-poverty neighborhoods experience improved mental and physical health [Sanbonmatsu et al., 2011]. We contribute to this debate by examining the differences in long-term labor market outcomes between public and voucher-subsidized housing.

The core identification challenge facing all research on subsidized housing is how to overcome the selection problem associated with a household's decision to participate in the program. That is, households that decide to participate in public housing and voucher-assisted housing are different from those that do not. The growing literature that uses instrumental variables procedures, experimental evidence, or quasi-experimental evidence regularly finds that the impact of subsidized housing is more positive when unobserved heterogeneity is taken into account.<sup>76</sup>

We make use of the large sample size and longitudinal nature of the administrative data set available to us and employ a household fixed-effects specification that exploits variation in voucher-supported housing and public housing participation across siblings over time within households. This allows us to isolate the effect of each type of subsidized housing on labor market outcomes from observed and unobserved household-level heterogeneity that may impact both labor market outcomes and the program participation decision.

Our results confirm that selection into subsidized housing matters. Whereas OLS estimates show a substantial negative effect of housing subsidies when young on later young adult earnings and employment outcomes, the household fixed-effects estimates are substantially less negative and, for many demographic groups, significantly positive. For example, for females, we find that being in public housing as a teenager yields a 29 percent premium for young adult earnings, and voucher housing, a 14 percent premium. These positive effects for females are mostly driven by the estimated effects for Black non-Hispanic households. Our approach, while superior to naïve OLS estimates, still may be subject to time-varying unobserved characteristics related to both adult earnings and household subsidized housing participation. We include several sensitivity checks to address these concerns.

The remainder of the paper proceeds as follows. Section 2 provides background information and reviews the existing literature. Section 3 describes the data and Section 4 the

---

<sup>76</sup>See the more detailed literature review below for relevant citations.

research design, hypotheses, and identification issues. Section 5 describes the study sample and Section 6 provides the empirical results. Section 7 concludes.

## **2 Literature Review**

Of the nearly 9 million Americans who participated in any subsidized housing program in the year 2000, well over half (6.4 million) were either in public or voucher-assisted housing. As of 2000, about 1.8 million American households lived in voucher-supported housing and about 1.3 million lived in public housing, made affordable by HUD subsidies. An even larger number were income eligible but elected or were unable to participate. In recognition of the fact that the primary types of federally subsidized rental housing differ along important dimensions, this paper compares the effects of living in subsidized rental public housing and subsidized voucher housing. In 2000, a central year in our study, 45 percent of public housing households and 61 percent of voucher households included children (see Table 1). A description of the major federal housing assistance programs that we consider appears in Appendix A.

HUD defines eligibility for its assistance programs based on family income as a percentage of Area Median Income (AMI), which adjusts for area income and for family size. Under most HUD programs, households pay 30 percent of their income for rent with HUD subsidizing the remainder to cover operating costs or up to a fixed local “Fair Market Rent”. Actual program requirements vary by subsidy type, but generally require residents to earn less than 80 percent of AMI (low income), with additional requirements dictating the percentage of residents that must be “very low income” (at or below 50 percent of AMI) or “extremely low income” (at or below 30 percent of AMI).

There is a broad literature estimating the economic effects of housing subsidies, although studies of the long-run impacts on children are scarce. In this literature, conclusions about the effects of subsidized housing vary considerably. In part, the mixed results

are likely a reflection of different study designs – many of the studies estimate the impact of moving from one type of subsidy to another. While certainly an interesting and policy relevant parameter, these studies are unable to answer how the different subsidy types compare to receiving no subsidy. Others that do compare subsidized households against non-subsidized households do not distinguish among different subsidy types and thus miss potentially important distinctions among the different programs.<sup>77</sup> Studies that have been able to compare multiple subsidy types to private, unassisted households typically do so for a limited geographic area (a city or metropolitan area) and focus on short-term, rather than longer-term outcomes.<sup>78</sup>

As we discussed in the introduction, the biggest challenge is dealing with selection. Subsidized housing residents differ observably from non-subsidized housing residents, often having characteristics typically associated with worse employment and educational outcomes. This implies that extensive controls are needed if the identification approach uses a selection on observables approach. However, if subsidized housing residents also have unobserved characteristics associated with worse labor market outcomes, then the estimated effects of assisted housing on outcomes are likely biased (for a general discussion, see Shroder, 2002). Previous research has employed a variety of approaches to deal with unobserved heterogeneity.

Some prior work relies on propensity score matching and other control variable-based methods to measure how outcomes differ among households in different public housing projects or programs. For example, Susin [2005] uses a rich set of controls from survey data to match households from project-based subsidized housing and Section 8 Housing

---

<sup>77</sup>For example, Olsen et al. [2005] used longitudinal HUD administrative data from 1995 to 2002 combined with data from other sources and a large, nationwide random sample to assess the employment results of multiple types of assistance. The authors found that each type of housing assistance has substantial negative effects on labor earnings that are somewhat smaller for tenant-based housing vouchers than for project-based assistance.

<sup>78</sup>For example, Bania et al. [2003] compares welfare leavers who received Section 8 housing vouchers or project-based housing, with other welfare leavers. The study was limited to Cuyahoga County (Cleveland), Ohio, and followed residents from 1996 through 1997 using administrative data. They found no significant effect from the receipt of housing assistance, and no difference between voucher and project-based assistance recipients.

with low-income non-recipient control households in the Survey of Income and Program Participation. He finds that housing subsidies reduce incentives to work and reduce earnings relative to control individuals but finds no difference in outcomes between voucher and project-based assistance recipients. However, he acknowledges some potential biases; for example, households with permanently low incomes may be matched to those with temporarily low incomes.

Recent work by Carlson et al. [2012b, 2012a] also uses a propensity score approach. As in this paper, the authors focus on employment and mobility outcomes for those receiving housing vouchers. The data come from administrative records in two databases maintained by the State of Wisconsin combined with Census Bureau public use microdata. Specifically, they draw a sample from the state's Client Assistance for Re-employment and Economic Support system, with 12,170 cases in the voucher group and 342,000 cases in the control group for up to 6 years after receipt of vouchers. Because the entire sample receives some sort of public assistance, their identification strategy is the equivalent of propensity score matching with a "receives public assistance" fixed effect. The results suggest that, 6 years after voucher receipt, there is little effect on employment, but housing voucher recipients experience a negative effect on earnings that diminishes over time [Carlson et al., 2012b]. Additional work indicates that voucher receipt resulted in both short- and long-term mobility and had little to no effect on four measures of neighborhood quality in the short term, but led to small long-term improvements in all four quality indicators [Carlson et al., 2012a].

The aforementioned studies rely on selection on observables identification approaches. With a rich set of controls, they are able to make progress on difficult selection issues but are subject to the concerns raised by Shroder [2002] and discussed above. As such, the literature has increasingly moved to alternative identification approaches to deal with these issues. Some researchers have used instrumental variables (IV) to identify the effect of public and assisted housing on outcomes. In one such study, Currie and Yelowitz [2000]

identify a regulation in housing assignment that provides an extra bedroom to households with two children of different genders, as compared to those with two children of the same gender. They use this rule based variation as an instrument to estimate the effect of public housing on child outcomes. They find that households entitled to an extra room because of the gender composition of their children are 24 percent more likely to participate in public housing and their children are less likely to have been held back in school. This suggests public housing participation has a positive impact on children's educational outcomes.

Newman and Harkness [2000] also use an IV strategy to identify the effect of participation in public housing on children's educational attainment. With a sample of about 1,000 individuals from the Panel Study of Income Dynamics, they develop a county-level measure of public housing availability by regressing the number of assisted housing units per income-eligible family in each county on county characteristics, and use the regression residuals for each county as an instrument for public housing participation. The authors find that public housing has no effect on children's education.

Other research takes advantage of public initiatives that resulted in the random assignment of program participation. For example, the Gautreaux project, which ended in 1998, involved the Chicago Housing Authority (CHA) distributing Section 8 housing vouchers to 7,100 African-American families on welfare. The vouchers were to be used to rent private market apartment units in either suburban or urban locations chosen at random by the CHA. Rosenbaum [1995] surveys 332 adults from the Gautreaux sample, and conducts detailed interviews with another 95. He finds that adult suburban movers (voucher recipients) experienced higher employment but no change in wages or hours worked relative to control adults.

Inspired in part by the Gautreaux project, HUD's Moving to Opportunity (MTO) project randomly assigned 4,600 households living in public housing projects in five cities to receive Section 8 housing vouchers, either with no restrictions or only for use in areas with a poverty rate below 10 percent. Despite the fact that MTO generated persistent improve-



ments in neighborhood conditions for treatment households, there was no significant effect on employment or earnings outcomes for adults or their grown children as reported by the parents [Sanbonmatsu et al., 2011].

The Welfare-to-Work Voucher Program provided housing vouchers to 50,000 families receiving or eligible to receive welfare. Mills et al. [2006] use an 8,371-household sample from seven public housing agencies to evaluate the differences in outcomes between those receiving vouchers and those not receiving vouchers. They find that vouchers somewhat improve the neighborhoods in which extremely low-income families live, but over a 3½-year study period, vouchers had no impact on employment or earnings.

Jacob [2004] makes use of the schedule of public housing demolitions in Chicago, and the Chicago Housing Authority's (CHA) policy of providing residents of demolished projects with Section 8 housing vouchers, to generate plausibly exogenous variation in public housing and Section 8 voucher participation. After matching administrative data from the Chicago Public Schools containing places of residence and test scores for 94,000 students to public housing addresses, Jacob finds that children leaving public housing fared no better or worse than their peers who remained in public housing for longer.

Jacob and Ludwig [2012] evaluate a CHA program that randomly assigned applicant households to a position on a waiting list for housing vouchers. Of the 82,607 households who applied for Section 8 vouchers between 1997 and 2003, they focus on the 90 percent of applicants living in private-market housing. Thus, they are able to compare housing voucher recipients to households who do not participate in subsidized housing. They find that vouchers reduce quarterly employment rates and earnings and increase participation in the Temporary Assistance for Needy Families program.

Oreopoulos [2003] uses another quasi-experiment, the random initial assignment of households to heterogeneous housing projects in Toronto, to estimate neighborhood effects on children. By matching earnings from Canadian tax data to historical information on parental residential location, he is able to reconstruct the childhood public housing experi-

ences of adult workers. The results indicate that neighborhood conditions as a child have no effect on adult earnings or welfare participation.

### 3 Research Design, Hypothesis, and Identification Issues

Our primary goal is to identify the causal effect of living in subsidized rental housing as a teenager on eventual labor market success. To do so, we begin by specifying a linear, constant effects regression model for a particular labor market outcome (the inverse hyperbolic sine of total earnings from 2008 to 2010 in this paper),  $y$ , of teenager  $i$  as

$$y_{if} = \alpha + \beta'H_i + \phi'X_{if} + \gamma'Z_{if} + \varepsilon_{if} \quad (47)$$

Where  $f$  indexes the household including child  $i$  in the year 2000. The outcome measures the teen's earnings as an adult;  $\alpha$  is an intercept. The variables of interest,  $H_i$ , are dummy variables that measure participation in subsidized housing (public housing or housing voucher) as a teenager. The vector  $X_{if}$  includes observable child and household control variables. The vector  $Z_{if}$  contains a set of unobserved characteristics that may be related to  $y_{if}$ . Lastly,  $\varepsilon_{if}$  is an independent error term.

Further, suppose that  $Z_{if}$  can be partitioned into two separate parts,  $Z_{if} = \begin{bmatrix} Z_f \\ Z_i \end{bmatrix}$ . Similarly,  $\gamma' = [\gamma'_f, \gamma'_i]$ . The first factor,  $Z_f$  is the composite of all observed and unobserved time-invariant characteristics for each household  $f$  that are common to all children  $i \in f$  and  $\gamma_f$  is the associated effect. The remaining factor,  $Z_i$ , contains other unobserved characteristics that vary by child, such as behavioral characteristics or disability status.

Consider estimating equation (47) using Ordinary Least Squares (OLS) and, thereby, omitting the unobserved characteristics in  $Z_{if}$ . The estimated coefficient  $\hat{\beta}_{OLS}$  will include both the true effect of subsidized housing participation and a term arising from omitted variable bias. The sign of the bias will depend on the effect of the omitted, household-

specific characteristics on earnings ( $\gamma$ ) and the covariance between participation in subsidized housing and the omitted characteristics. For example, if households with unobserved characteristics that tend to depress child outcomes are also more likely to enter public housing, then  $\hat{\beta}_{OLS}$  will be biased downward. Thus, a finding that subsidized housing depresses child outcomes may be spurious unless the specification controls for these potential biases. To account for the possibility that estimates are contaminated by household-level heterogeneity, we propose an alternative identification strategy. To the extent that the bias in OLS estimates is solely attributable to the omission of time-invariant heterogeneity at the household level that is correlated with both program participation and labor market outcomes, then conditioning on household fixed effects will eliminate the bias.

To that end, we specify a household fixed-effects regression that exploits within-household variation in program participation (across siblings) to identify the impact of housing subsidies.

Griliches [1979] provides a summary of the early literature that makes use of sibling fixed effects and points out a number of potential issues. Recent studies include Royer [2009] who used over 3,000 twin pairs and twin fixed effects to estimate the effect of birth weight on long-term outcomes, Currie and Walker [2011], who used mother fixed effects to estimate the impact of the introduction of EZ-Pass in New Jersey and Pennsylvania on infant health outcomes, and Currie et al. [2010], who employed sibling fixed effects to identify the relationship between early childhood health problems and outcomes in early adulthood. An especially relevant siblings study is Aaronson [1998], who estimated the effect of neighborhood characteristics on children's educational outcomes. Aaronson used the Panel Study of Income Dynamics to examine over 2,000 individuals in over 700 families and measures differences in exposure to high poverty neighborhoods across siblings. He found negative effects on high school graduation with and without the household fixed effects.

In our study, the household fixed-effects estimates control for time-constant, unob-

served household-level heterogeneity ( $Z_f$ ). The household fixed-effects (HFE) regression estimates the effect of subsidized housing participation on labor market outcomes using only variation in housing participation and outcomes across teenagers within the same household. In practice, we subtract out the household mean of each dependent and independent variable from each observation within a household.<sup>79</sup> Therefore, HFE only uses observations from household  $f$  to help identify  $\hat{\beta}_{HFE}$  if there are at least two teenagers  $i$  and  $j$  aged 13-18 in the household in 2000 where  $H_i \neq H_j$ . For example, consider a household in the year 2000 with a 17 year-old and a 14 year-old that does not enter HUD-subsidized housing until 2003. The older sibling, who leaves the household in 2002, would have  $H_i = 0$  and the younger sibling would have  $H_i = 1$ .

The HFE model is written as:

$$y_{if} = \alpha + \beta'_{HFE} H_i + \phi' X_i + \gamma_f + \gamma' Z_i + \varepsilon_{if} \quad (48)$$

Where  $\gamma_f$  gives the fixed effect for all children in household  $f$ . The effects of observed characteristics common among all children in a household are not separately identified, but instead subsumed in  $\gamma_f$ , so only a subset of  $X_{if}$  remains. In practice,  $H_i$  is a vector containing measures of participation in both public housing and housing voucher programs as a teenager,  $X_i$  contains an indicator for whether the teenager is male, a set of age dummies, and, in some specifications, an interaction between whether the teenager is male, the full set of age dummies, and the race/ethnicity of the household. We also interact each of the subsidized housing measures with whether the teenager is male to allow for heterogeneous effects by teenager gender, and we estimate separate regressions for each race/ethnicity to allow all coefficients to vary along this dimension. We estimate both a “dummy” version where the “treatment”  $H$  is a set of two binary indicator variables for whether an individual resided in each type of subsidized housing as a teenager and a “dose” version where treatment is the number of years an individual resided in each type of subsidized housing

---

<sup>79</sup>We also cluster standard errors at the household level.

between ages 13 and 18.

The HFE estimation provides an unbiased estimate of the effect of teenage subsidized housing residence on labor market outcomes under much less stringent conditions than a typical conditional on observables approach (including propensity-score matching approaches which also hinge on controlling for all relevant observables that determine selection and impact outcomes). There are, however, two types of characteristics contained in the child-specific factor,  $Z_i$ , that could lead to bias in  $\hat{\beta}_{HFE}$ . First, any household-specific and time-varying characteristic that is correlated with both subsidized housing residence and labor market outcomes will lead to bias. For example, if families enter subsidized housing in response to negative economic shocks, and these negative shocks are also harmful to the subsequent labor market outcomes of the child,  $\hat{\beta}_{HFE}$  would be a downward-biased estimate of the true effect.<sup>80</sup> In fact, HUD strongly prefers and in some cases requires that program households be below a certain income threshold. This suggests that if any bias from unobserved, time-changing heterogeneity is present, this bias is likely to be negative. To address this possibility, we also consider HFE specifications where we control for the parents' earnings while the teenager is between 13 and 18. This variable will capture differences in parental earnings across siblings that have different subsidized housing experiences.<sup>81</sup>

A second potentially confounding unobserved characteristic is any within-household, teenager-level heterogeneity that is correlated with both labor market outcomes and subsidized housing participation. In this case, the direction of the potential bias is less clear. However, it seems implausible that this type of bias would contaminate the HFE estimates.

---

<sup>80</sup>Job loss by a household member is an example of an economic shock, though it is unlikely that housing subsidies are responsive to transitory events as the waiting lists are typically substantial. Another plausible scenario given eligibility requirements imposed by HUD is that households are more likely to be admitted into subsidized housing after a household member develops a disability. Again, under the assumption that exposure to this disability worsens potential labor market outcomes, this would lead to a downward-biased estimate.

<sup>81</sup>Aaronson [1998] also evaluated the validity of using across-sibling variation by examining whether moves into or out of high-poverty neighborhoods co-vary with other household characteristics, such as parents' income.

The decision to move into subsidized housing is made at the household level. In effect, for this to be a concern, households would have to be making housing decisions in response to the characteristics of one teen but not the characteristics of the other teenage household members. Another factor that might mitigate concerns of correlation of housing treatment and child characteristics is the waiting periods typical for receipt of a housing subsidy. Such delays would tend to reduce any correlation of housing treatment and unobserved characteristics, which should attenuate any bias. Indeed, waiting times are one reason that siblings may have different housing treatment experiences (in terms of dummy or dose), which is ideal variation for our analysis.

## 4 Description of the Data

### Siblings Sample Frame

The core data set brings together person- and household-level records from the 2000 Decennial Census and several different administrative files. To begin, we use the responses from the 2000 Census to construct a frame of over 1.8 million youth aged 13-18 and their households. Because our focus is on employment outcomes from 2008 to 2010,<sup>82</sup> we require that children are at least age 13 in 2000, meaning they will be at least 21 by 2008 and may be entering the labor force even if they attained some higher education. We cap the sample at age 18 and require that in 2000 the child was in a household with their parent(s). Including older youths would undermine the focus of the paper, and our identification approach relies on the assumption of parents making housing decisions for children.

Because our aim is to estimate the effect of childhood environmental factors on later life outcomes, we derive most of our demographic characteristics from the base year 2000 Census short form responses, when subjects are still children.<sup>83</sup> We retain responses for

---

<sup>82</sup>We recognize that 2008-2010 is a sluggish period for the national labor market, but our identification approaches are designed to exploit the cross-sectional variation. In future work we may consider whether the effects vary across the business cycle.

<sup>83</sup>We chose to use all households in the U.S. rather than the 1-in-6 sample filling out the long form

one or two parents as well as all youth between the ages of 13 and 18.<sup>84</sup> We use time-invariant explanatory variables relating to the child such as date of birth, gender, race, and ethnicity, and characteristics of the household in the base year such as housing tenure (rent or own), number of people, number of children.<sup>85</sup> We also construct a household race/ethnicity variable to allocate households to race/ethnicity subsamples. Specifically, we define a household as Hispanic if any member reports being Hispanic, Black non-Hispanic (Black) if no member reports being Hispanic and at least one member reports being Black, White non-Hispanic (White) if no member reports being Hispanic or Black and at least one member reports being White, and Other non-Hispanic (Other) if no member reports being Hispanic, Black, or White.

Youth in the Census 2000 frame are then matched to administrative records on housing subsidies from the Department of Housing and Urban Development's HUD-PIC<sup>86</sup> file, place of residence from the Longitudinal Employer-Household Dynamics (LEHD) maintained Composite Person Record (CPR), and subsequent earnings from the LEHD<sup>87</sup> using a unique person identifier. Person-level record matching is done by way of a Protected Identification Key (PIK), which is assigned to survey and administrative records based on personally identifying information. The 2000 Census has PIKs for over 89 percent of the person-records, while almost 98 percent of HUD records have a PIK, and all LEHD records have a PIK. We only retain households with a parent who has a PIK and at least two children aged 13 to 18 that have a PIK and non-missing basic characteristics.<sup>88</sup> From the full sam-

---

for the principal analysis in order to have a larger sample size. While the long form would allow us to include variables such as parent's education, such time-invariant explanatory factors will be subsumed into the household fixed effects in any case.

<sup>84</sup>We define the head of household and the spouse of the head of household as the parents for each MAFID. MAFIDs, or addresses, constitute the residence frame for Census Bureau surveys. In some cases these individuals may be grandparents, other relatives, or even unrelated adults.

<sup>85</sup>We exclude households including more than 15 residents or more than 10 teenagers.

<sup>86</sup>PIC refers to Public and Indian Housing Information Center. The data file contains an annual extract of recipients of voucher-supported housing and public housing, submitted by housing authorities and providers. For other research using the HUD-PIC extract file, see Lubell et al. [2003]; Mills et al. [2006]; Olsen et al. [2005]; Shroder [2002]; and Tatian and Snow [2005]. We do not use the HUD-TRACS (Tenant Rental Assistance Certification System) since those data apply to project-based Section 8 subsidies.

<sup>87</sup>For a description of the LEHD infrastructure files and public statistics, see Abowd et al. [2004].

<sup>88</sup>For cases where a PIK has been assigned to multiple responses (less than 1 percent) we drop all cases,

ple including records with no PIK, we estimate a logistic regression for whether or not a person response has a PIK, with explanatory variables including the number of persons in a household, the number of children, housing tenure as well as person age, gender, race, ethnicity and state fixed effects based on the year 2000 location.<sup>89</sup> To retain a representative sample of records with a PIK, we reweight them using the inverse of the probability of having a PIK, based on the model.

### **Housing Subsidy**

The HUD-PIC file provides detailed information on public housing and Housing Choice Voucher recipients during our study period from 1997 to 2005. As part of their housing occupancy verification process, local housing authorities provide HUD with the identities of residents, which HUD then compiles into an annual relational database. Table 1 presents characteristics of public and voucher supported housing participants from public use data derived from the HUD-PIC file. In 2000, households averaged approximately \$10,000 in annual income, which was about a quarter of metropolitan area median income.

The person-level file used at the Census Bureau includes demographic and housing unit information, but this study only makes use of occupancy as an indicator of housing treatment. We match PIKs from the decennial file to the HUD-PIC file and identify whether a child resided in public or voucher housing in each year from 1997 to 2005. We consider a child to be a HUD-subsidized resident in a particular year if their PIK appears in the HUD administrative data and if that individual is still under the age of 18.<sup>90</sup> Thus, the maximum number of years a child could reside in HUD housing is 6 years before turning 18, which could occur for a 13-year-old first residing in subsidized housing in or before 2000. An 18-year-old in 2000 could only reside in HUD-subsidized housing for at most 4

---

unless all observable characteristics (date of birth, race, ethnicity, gender, geographic location) are identical, in which case one record is retained.

<sup>89</sup>Characteristics highly associated with not having a PIK include race, ethnicity, age, and sex.

<sup>90</sup>We do not count individuals who are under 18 in 2000 but over 18 when we observe them in the HUD administrative data as being HUD residents.



years (beginning in 1997).

We construct an indicator variable for whether a teen resided in either public or voucher housing any time between 1997 and 2005. Our goal is to estimate the effect of this binary treatment variable on labor market outcomes. We also examine the effect of a treatment “dose” variable that could take on values from 0 to 6 for the count of (post-1996) years a child resides in voucher or public housing.

There are three exclusions we found necessary to avoid sample contamination due to possible measurement error. First, we exclude all households owning a home in 2000 (based on the decennial census response). While individuals in such households could end up in subsidized housing later in the decade, we decided they may not be representative of renter households eligible for subsidized housing. Second, we require that each teenager’s parents earn less than 50 percent of HUD-specified Area Median Income (AMI) on average while the child is a teenager.<sup>91</sup> Third, we excluded households who lived in the 119 counties participating in HUD’s Moving to Work (MTW) demonstration (see Abravanel et al. 2004). Local housing authorities participating in the demonstration were permitted to stop reporting administrative data to HUD on participants. We excluded these counties to avoid misinterpreting changes in housing authority reporting requirements as within-household differences in subsidized housing participation.

### **Labor Market Outcomes**

LEHD, a partnership between the Census Bureau and all 50 states and the District of Columbia, produces public use data tabulations that are widely used by state and local governments.<sup>92</sup> At its core are two administrative records files provided by states on a quarterly basis: (1) unemployment insurance (UI) wage records, giving the earnings of

---

<sup>91</sup>We use average annual total labor income from years where the child is between 13 and 18 years of age. To avoid dropping observations that do not match to the Composite Person Record (CPR) we use the 2000 census residence county to define AMI. HUD defines AMI using American Community Survey data; specified proportions of AMI are used as eligibility and priority criteria.

<sup>92</sup>LEHD data products include the Quarterly Workforce Indicators, the LEHD Origin-Destination Employment Statistics (LODES), Job-to-Job Flows.

each worker at each employer, and (2) employer reports giving establishment-level data, also known as the Quarterly Census of Employment and Wages (QCEW), but often referred to as the ‘ES-202’ program. The coverage is roughly 96 percent of private non-farm wage and salary employment [Stevens, 2007].<sup>93</sup>

The longitudinal LEHD data are based on quarterly earnings information for more than 130 million U.S. workers and their employers covered under state UI systems beginning in the mid-1990s and continuing to the present, essentially a universe of workers. The longitudinal data thus permit the measurement of complete employment “histories” beginning with a person’s entrance into the labor force. This information includes earnings, employment status and industry, along with other work and home location information. Thus, LEHD wage data matched to the Census 2000 data enable us to track a large set of children into adulthood and measure earnings and employment outcomes. For our purposes, the national nature of the files and complete work histories enable us to compute outcome measures for individuals over any given horizon such as the number of quarters worked, cumulative number of jobs, the number of spells of joblessness, the durations of spells of joblessness, and the earnings levels and its growth within and between jobs.

For regression purposes, we use the inverse hyperbolic sine (IHS) of earnings<sup>94</sup> rather than the more traditional log of earnings because estimated coefficients can be interpreted in the same way as with a log transformed dependent variable but, unlike with the log of earnings, IHS is defined for zero earnings. The IHS is defined as  $y_i^* = \log \left[ y_i + (y_i^2 + 1)^{\frac{1}{2}} \right]$  where  $y_i$  is total earnings for individual  $i$  (see Burbidge et al., 1988).

---

<sup>93</sup>LEHD is in the process of integrating data on self-employed individuals and independent contractors who are not covered in the UI files but are available from the Census Bureau’s Business Register which contains the universe of all businesses including all sole proprietorships on an annual basis (whether the sole proprietor has employees or is a non-employer). In addition, the LEHD project has acquired the personnel records from Office of Personnel Management (OPM) so that federal workers are now also tracked in the file system. This study does not yet make use of these new data sources, but may in future versions. For more information on the LEHD, see Abowd et al. [2004].

<sup>94</sup>Annual earnings are deflated to their 2000 purchasing power equivalent using the U.S. city average annual purchasing power for all urban consumers.

## Other Factors Varying Within Households

We introduce additional geographic data to address time-varying but spatially constant household factors. The LEHD program maintains an annual place of residence file composed of federal administrative data known as the Composite Person Record (CPR). LEHD uses CPR residences, which begin in 1999, for imputation models and for the residence component of public use data. We identify a residence census block for each child from 1999-2005 where available (approximately 10 percent of children are missing a CPR residence in each year). Where possible, we match the child residence to block group-level tabulations from Census 2000, giving neighborhood characteristics such as the poverty rate.

In addition to using LEHD earnings to construct outcome measures for the teenagers, we use parent's LEHD earnings to determine sample eligibility and to construct an annual measure of household income for 1997 to 2005 to use as a control variable. HUD defines eligibility for its assistance programs based on family income as a percentage of Area Median Income (AMI), which adjusts for area income and for family size.<sup>95</sup> For each teen, we calculate average parents' earnings (the sum of earnings for the head of household and the spouse of the head of household) while the child was a teenager (also transformed into the IHS of average income to match the dependent variable). Additionally, we use each household's location in 2000 and household size in 2000 matched to their average parents' LEHD earnings to identify AMI figures at the county level. We then create a ratio of parents' earnings to AMI in order to account for the differences in average earnings across regions, which can vary by almost \$75,000 for metropolitan areas within the U.S. Since local housing authorities often require that a household earn less than 50 percent of AMI to be eligible for assistance, we retain only children in households with a parents' earnings-to-AMI measure below 0.5, so that the analysis sample includes only those widely eligible

---

<sup>95</sup>Under most HUD programs, households pay 30 percent of their income for rent with HUD subsidizing the remainder or up to a fixed local "Fair Market Rent." Actual program requirements vary by subsidy type, but generally require residents to earn less than 80 percent of AMI (low income), with additional requirements dictating the percentage of residents that must be "very low income" (at or below 50 percent of AMI) or "extremely low income" (at or below 30 percent of AMI).

for the subsidized housing treatment. As with the labor market outcomes, some households may appear to have lower incomes because they do not work in UI-covered employment. In future work, we will assess the significance of such omissions for our sample composition.

We employ both the composite of neighborhood (at the Census block group-level) poverty and the IHS or average annual parents' earnings between the ages of 13-18 as control variables in some specifications. Aaronson [1998] examined whether cross-sibling variation in household income is associated with moves across neighborhoods. Likewise, we acknowledge that changes in household income may be directly associated with moves into and out of subsidized housing. Controlling for the household income during the period each sibling is between 13-18 should alleviate these concerns. Controlling for changes in the poverty rate when each sibling is between 13-18 is designed to capture one of the mechanisms for the impact of subsidized housing. As such, we interpret adding each of these two longitudinal controls somewhat differently. We interpret specifications with controls for parents' earnings as a robustness check on the importance of unobserved, time-varying characteristics, and those with controls for block group percent poverty as a test of one potential causal mechanism.

## **5 The Sample: Basic Facts**

In sum, to be included in the estimation sample, we require that individuals have been between 13 and 18 years of age in the year 2000, have non-missing values for age, gender, ethnicity, treatment status, and residential location, have successfully been assigned a unique PIK based on the 2000 Census, and be from the same 2000 renter household as at least one other teenager. Finally, because not all households are eligible for subsidized housing, we limit our sample to teenagers from households more likely to qualify for housing assistance, with average annual earnings below 50 percent of local AMI (see above). Of the 2.8 million children aged 13-18 in the U.S. in 2000, we end up with a final sample

size of 521,000 teenagers.

Table 2 presents summary statistics for this sample.<sup>96</sup> The first column presents summary statistics for the sample used in estimation – teenagers living with another sibling aged 13-18. This sample is subdivided further, into those who were in households not in subsidized housing anytime during the 1997-2005 study period (column 2), and those who were (column 3); the latter are then subdivided further, into those who never lived in subsidized housing as a teenager 13-18 (column 4), and those who did (column 5). The comparison between columns 4 and 5 is the raw differences analog to our main empirical results.

There are a few minor differences between the estimation sample (column 1) and the full sample of teenagers; that is, the sample including cases in which there is only one teenager in the household (not shown). Of course, since we require that the estimation sample have at least two teenagers aged 13-18, the average household size is bigger. In the estimation sample, the proportion which is non-Hispanic Black is slightly higher, the proportion in single-parent households is slightly lower, and the proportion receiving a housing subsidy is slightly higher. These differences relate to the generalizability of the study, but have no bearing on the identification approach.

Comparison of columns 2 and 3 foreshadow the likely findings from an OLS regression. There are substantial differences in the outcome variables examined – those in subsidized housing earned less during the 2008-2010 period (\$24,651 versus \$32,443 on average), they worked fewer quarters (6.540 versus 7.209 on average), and a lower percentage had any labor market earnings during the 2008-2010 period (80.6 percent versus 83.3 percent). Comparisons of columns 4 and 5 foreshadow the likely findings from a household fixed effects regression—there are few differences apparent from the comparison. But unconditional differences are not likely to tell the whole story, nor do these comparisons only use

---

<sup>96</sup>Confidentiality restrictions preclude us from releasing summary statistics for the entire sample of 13-18 year old children from the 2000 census.

within-household variation, so we now turn to the regression analysis.<sup>97</sup>

Figure 1 displays the distribution of within-household differences—each teenager’s own subsidized housing participation net the household mean for all teenagers—we use to identify our regression model. The figure is based on the sample in Table 2, Column 3, but teenagers are also required to be from households with at least some within-household difference in subsidized housing participation among the household members aged 13-18.<sup>98</sup> This sub-sample included 41.7 percent of housing voucher participants and 69.3 percent of public housing participants. The distribution is unimodal and symmetric around zero, with an overwhelming majority of teenagers within two years of the household mean participation.

Given the identification strategy we employ, an important question is what causes the observed within-household differences in subsidized housing participation?

We define treatment as teenage—i.e. between the ages of 13 and 18—participation in subsidized housing. However, for sample members who are 17 or 18 years of age in 2000, we are unable to observe their subsidized housing participation at age 13 (or age 14 for individuals aged 18 in 2000) because our administrative records begin in 1997. As a result, it is possible that some of the within-household variation results from this left-censoring of treatment. We test for the importance of censoring by limiting the sample to only teenagers aged 13-16 in 2000; that is, those teenagers with uncensored treatment. We find no differences between our main estimates and the results run on the age-limited sample.<sup>99</sup> Therefore, while some of the observed within-household variation may result from age censoring, this variation does not drive the empirical results we present later.

Similarly, measurement error in the administrative subsidized housing records could create within-household variation. For example, if for some reason HUD’s enumeration of

---

<sup>97</sup>Only 15 percent of children in the ever-subsidized household sample receive no subsidy between the ages 13-18. This might seem to be a small subset to serve as a “control” sample for the effect of a subsidy in the dummy treatment variable regressions. Note, however, that we also estimate models with a dose treatment variable, allowing for wider variation in subsidy receipt.

<sup>98</sup>The restriction that teenagers have some within-household variation is made for expositional purposes.

<sup>99</sup>Results available upon request.

children in a household is incomplete in one year, we would incorrectly interpret the incomplete record as within-household heterogeneity in subsidized housing participation. To account for this possibility, we present a robustness check that predicts subsidized housing participation for each teenager using just their age and the observed participation of their head of household from the 2000 census. This predicted treatment is not subject to differential measurement error within a household. We show both reduced form estimates that use predicted participation to define treatment and IV estimates that instrument for actual participation with the predicted participation. The results suggest that measurement error in the administrative records does not drive our estimates or explain the within-household differences we observe.

A third possibility, is that changes in parental income or earnings could alter household eligibility for different types of subsidized housing. This is potentially problematic for our identification strategy as household fixed effects do not account for this type of time-varying heterogeneity. As we discuss in more detail later in the paper, we develop a longitudinal, child-specific measure of parental earnings using the LEHD data and test whether our results are affected by its inclusion as a control variable. In short, the main estimates are unaffected. Given the extensive literature suggesting parents earnings have a strong positive relationship with child earnings (see Chetty et al., 2014 for a recent example), this suggests that within-household (i.e. longitudinal) changes in parental income and earnings are unlikely to explain much of the within-household variation in subsidized housing participation.

In their research estimating the impact of the Earned Income Tax Credit (EITC) on earnings, Chetty et al. [2013] use changes in information about the existence and shape of the EITC benefit for identification. As with EITC, it is possible that eligible households are not aware of the location of public housing projects, their own eligibility for public or voucher-assisted housing, or how to apply to either program. If households acquire information about one or both programs while they have multiple teenage children,

it could prompt them to apply for subsidized housing or to switch between the two programs. This would then create within-household variation in teenage subsidized housing participation. Further, this within-household variation would be unrelated to the potential outcomes of children, having been driven instead by the timing of an information shock to the household. Unfortunately, while this seems plausible and remains a likely explanation for within-household variation in treatment, we have no way to test empirically for the existence of such information shocks.

Finally, as pointed out by Jacob and Ludwig [2012] and others, subsidized housing programs are frequently oversubscribed, leading to lengthy lags between when households apply for a particular program and when they are allotted a voucher or public housing unit. Households that apply to an oversubscribed subsidized housing program may end up with children exposed to different amounts of the program purely as a result of their mandated wait time. Consider a household with one 13-year-old and one 12-year-old that applies for a public housing program, is placed on the waitlist for one year, and then remains in that project. In the absence of the wait time, both children would experience the same amount of teenage public housing participation: six years each. However, because of the one year wait, the 13-year-old will end up spending only five teenage years in public housing while the 12-year-old will spend six.

There also appear to be substantial wait times for both public and voucher-assisted housing in our sample. To illustrate these wait times, we use data on all subsidized housing participants from the year 2000. For most households, the data contain information on the date they entered a waitlist as well as the date they were granted admission to the program. In some cases the two dates are the same, indicating there was no wait for the program. For most households there was a non-trivial wait between when they were placed on a waitlist and when they were admitted. Figure 2 displays the distribution of wait times for individuals in voucher and public housing who entered subsidized housing no earlier than 1995 and who were found in subsidized housing in 2000. We restrict the entrance date to



be after 1995 because data quality decreases rapidly in the early 1990s and because these waits are likely to be a better approximation to the waits experienced by the households in our sample. Figure 2 indicates that about 12% of public housing residents and 29% of housing voucher recipients faced wait times of one year or more. Clearly, many perspective subsidized housing participants face lengthy lags between when they apply and when they are admitted to programs. These lags offer another plausible explanation for the observed within-household differences in subsidized housing participation.

## 6 Empirical Results

### Samples and Specifications

The key question we address is whether living in voucher-supported or public housing affects a teenager's labor market experiences as an adult. We compare the effects on earnings over the 2008-10 period of each of these two HUD housing types with nonsubsidized housing.<sup>100</sup>

Table 3 presents results for teenagers from all households while Tables 4, 5, and 6 present results for teenagers from non-Hispanic White households, non-Hispanic Black households, and Hispanic households, respectively. Each table presents results for a “dummy treatment,” which consists of a binary measure of whether an individual ever participated in each type of subsidized housing as a teenager, and a “dose treatment,” which is defined as the number of years an individual participated in each type of subsidized housing while under the age of 18. As described above, the dependent variable is the inverse hyperbolic sine of total earnings over the 2008-10 period. In addition to the treatment variables interacted with gender, unlisted controls include age, gender, age by gender, and

---

<sup>100</sup>In unreported results, we have also used the total number of quarters worked over the 2008 to 2010 period and an indicator for whether the individual ever worked during the 2008 to 2010 period as dependent variables. The results are qualitatively consistent regardless of which measure of labor market performance is used.

household race/ethnicity by gender.<sup>101</sup> Table 7 presents the effect of each type of housing subsidy, separately for each sex and household race-ethnicity type, and it compares the estimated effect across gender and across the two subsidized housing types within each possible sex/household race-ethnicity combination.

In Tables 3 through 6, the first column presents OLS estimates of the specification described in equation (47). The coefficients capture the correlation between earnings and the two different types of subsidized housing participation after controlling for observed covariates, but, as discussed before, are susceptible to bias as a result of selection based on unobservable factors. The second column in each table presents estimates from the household fixed effects (HFE) specification, described in equation (48). By using only within-household variation, these estimates purge the treatment effects of all bias resulting from time-invariant, household-level unobserved characteristics. We believe these estimates better capture the causal effect of subsidized housing participation as a teenager on adult labor market earnings.

The third, fourth, and fifth columns in each table presents results from a HFE specification that, in addition to the controls in column (2), also includes, in column (3), a control for the average parents' earnings that each individual experienced between 13 and 18 and its interaction with a male dummy, in column (4) a control for average block group percent poverty that each child experienced between 13 and 18 years of age and its interaction with a male dummy, and in column (5) controls for both parents' earnings and block group poverty and their interactions with male dummies. We interpret the estimates in Column 3 as a test for whether our household fixed effects are effectively ridding the treatment effects of bias from unobserved, time-varying heterogeneity. Specifically, if our treatment effects do not change after the inclusion of parents' earnings, then either the within-household differences in subsidized housing participation or the within-household differences in adult earnings (or both) are unrelated to within-household differences in parents' earnings. Sim-

---

<sup>101</sup>The complete regression results as well as the results for the other measures of labor market performance (cf. footnote 24) are available from the authors.

ilarly, the estimates in column 4 are an indicator of whether neighborhood quality, as proxied by block-group percent poverty, is a potential mechanism for the estimated treatment effects. Column 5 accounts for both factors.

### **Results for all Households**

We now turn to the coefficients of interest beginning with the estimates that pool across household race/ethnicity in Table 3. In column (1) , for both the dummy and dose treatments, the OLS results show that there are significant negative effects on subsequent total earnings with larger negative effects for males. Significant negative relationships between the two types of subsidized housing participation and adult earnings also occur in each of the race/ethnicity groups (Tables 4-6), although magnitudes vary.

However, the HFE results, which control for all household level time-invariant heterogeneity, paint an entirely different picture; the HFE results for females and males are summarized in Table 7, Panel A. The negative effects from OLS are attenuated or reversed. Housing voucher participation is not negatively related to adult earnings for females in the HFE specification. Both living in public housing and living in a housing voucher-subsidized unit lead to positive and significant effects on later earnings for females. The effect of voucher participation remains negative and statistically significant for males with the dummy treatment, and is not statistically different from zero for public housing. The effects estimated for the dose treatment (years) reinforces the findings of the dummy treatment. The effects for males are significantly more negative than the effects for females in both the dummy and dose case. For the dummy treatment, public housing is more beneficial than housing vouchers for both females and males (no difference was found for the dose treatment).

The point estimate suggests that ever having lived in voucher-supported housing as a teenager increases early adult earnings by roughly 14 percent for females and reduces earnings by roughly 24 percent for males. The dose results indicate that each additional

year of voucher participation increases adult earnings for females by about 6 percent and reduces adult earnings for males by 3 percent. For public housing, the relationship between participation and future earnings is not significantly different from zero for males, in both the dummy and dose cases, but is positive and significant for females. The point estimate suggests that for females, ever having lived in public housing as a teenager increases adult earnings by roughly 29 percent. The dose result suggest that each additional year of public housing participation increases early adult earnings for females by 9 percent.

The results in columns 3 and 4, which add controls for average parents' earnings and average block group percent poverty, are essentially unchanged. In the following subsection, we find that columns 2 and 3 are similar even when allowing the results to differ for different race/ethnicity samples. We believe this indicates that the household fixed effects specification is effectively ridding the treatment effects of bias from unobserved, time-varying heterogeneity. Consequently, we report just the simple HFE results in the text (and Table 7) below.

### **Race/Ethnicity Samples**

To help understand the results in Table 3, we investigate whether the results differ by household race-ethnicity. Tables 4 through 6 thus explore whether there is treatment effect heterogeneity by household race-ethnicity. We do this by estimating coefficients separately for non-Hispanic White households, non-Hispanic Black households, and Hispanic households. Comparing results across these three subgroups (see Table 7 for a summary), we find important differences. For example, the HFE results show substantial positive effects of both living in voucher-supported housing and in public housing on young adult earnings for Black females, but not for Black males or for any subset of Hispanics or non-Hispanic Whites (there is actually a negative effect of housing vouchers on earnings for Hispanic males). The dose results generally echo the dummy results, except that we find a positive effect on later earnings for non-Hispanic Black males who lived in public housing. Clearly,

there is important heterogeneity across race/ethnicity groups, affirming the importance of considering these groups separately.

The positive effects for non-Hispanic Black females suggest they receive an earnings premium of 15 percent from participating in the housing voucher program and 18 percent from living in public housing relative to not having participated in either program. The dose results indicate that each program increases earnings relative to non-participants by about 6 percent per year. Non-Hispanic Black males also see their adult earnings increase as a result of public housing participation, by about 7 percent per year of residence. The estimate for Hispanic males indicates that voucher housing decreases adult earnings by about 8 percent.

Table 7, in addition to displaying the average partial effects of each type of subsidized housing separately by gender, also displays tests of whether the effects of each type of subsidized housing are equal. For example, we test whether the effect of voucher housing for females is the same as the effect of public housing for females. We conduct this test for each possible household race/sex combination, and for both the dummy and dose treatments. For the combined sample, we find that vouchers lead to lower male outcomes than public housing for the dummy but not the dose treatments. For the subsamples, this result is apparently driven by the result for non-Hispanic Black and Hispanic males; there was one significant dose treatment difference – for non-Hispanic Black females. When comparing the estimates for females to those for males, the results differ for non-Hispanic Blacks in voucher housing, and for non-Hispanic Whites in public housing. Though there are differences between males and females for the combined sample, there are no statistically different effects when the samples are disaggregated by race/ethnicity.

## 7 Extensions and Robustness Checks

We undertake three extensions in an attempt to understand these results in more detail. As much of the discussion of public housing in the popular media concerns high-rise projects primarily found in urban areas, we check whether the effect of living in a large public housing project is different from the overall results. That is, we allow for the effect of public housing participation to differ according to project size (population). To do so, we defined person-weighted project size quartiles by considering all public housing projects over the period 1997-2005. On the basis of these quartiles, it was determined whether each individual in our sample who ever participated in public housing was also a resident of large public housing (the top quartile). We then included either an indicator for whether each teenager in our sample ever lived in a large public housing project or a count of the number of years each teenager lived in a large public housing project in addition to the measures of housing voucher participation and general public housing participation included in previous specifications. The coefficient estimates from household fixed effects specifications for these large public housing measures capture any differential effect that large public housing residence as a teenager has on adult earnings. Table 8 presents these results. The estimated coefficients on the housing voucher and general public housing measures are almost identical to those from the more basic household fixed effects specification. This suggests the heterogeneity with respect to project size in the effect of public housing is not particularly important empirically. In no column is the differential effect of large public housing significantly different from zero for females or for males. We therefore find no evidence to support the idea that living in a large public housing project is particularly harmful for children's later earnings.

Similarly, it might be the case that being assigned to a public housing project where households earn relatively low annual incomes has a differential impact on adult outcomes. Such a differential effect could exist as a result of role model effects (e.g. observing adults who supply more labor while a teenager increases labor supply as an adult) or if project

level social networks enable individuals to find a job or a higher paying job more easily. To test for heterogeneity by project-level household income, we compute the person-weighted median household adjusted income for each project year.<sup>102</sup> Next, we create year-specific quartiles and assign each project-year to a quartile. Teenagers in our sample are then matched to the public housing project and the associated household income quartile for each year they participated in public housing. We define the lowest-income public projects as those that fall into the bottom quartile with respect to median household annual adjusted income. This match is used to create an indicator for whether each teenager ever resided in a lowest-income public housing project and a count of the number of years they resided there. These measures are then included, in addition to the housing voucher and general public housing measures, as discussed in the previous paragraph. Table 9 presents the household fixed effects estimates from these specifications. Allowing for the effect of public housing to differ by median household income has almost no impact on the main estimates. Further, there appears to be no additional impact of living in one of the lowest-income public housing projects for either males or females. Taken together, Tables 8 and 9 indicate little heterogeneity in the estimated treatment effect of public housing along project-type dimensions.

One additional robustness check we conduct deals with possible measurement error as well as endogenous changes in the structure of households. Specifically, in Table 10 we use a predicted rather than the actual indicator of being in public or voucher housing. To construct the predicted value, we use the age of the children in the household and the actual information on whether the parent is in subsidized housing. That is, for any given year, if a parent is in subsidized housing and the child is in the 13-18 year-old age range, then the “predicted” participation measure indicates that the child is in subsidized housing. Differences between actual and predicted measures of participation might arise for two reasons, both of which we would like to avoid. The first is measurement error. The second is that

---

<sup>102</sup>HUD computes adjusted annual income on the basis of household-type (elderly, disabled, family), the number of dependents in the household and income net of certain child care, medical and disability expenses.

the child left the household while still aged 13-18. Such departures might reflect events (e.g. a child leaving to live with a member of the extended family such as a grandparent) that have an impact on later outcomes but are unrelated to the mechanisms we are seeking to identify. Using this approach, Table 10 reports results for using the actual treatment (the same as Table 3), using the predicted treatment instead of the actual treatment, and instrumenting the actual treatment with the predicted treatment. The results in Table 10 are strikingly similar when using any of the participation definitions. This suggests that measurement error and household departures/dissolutions are not importantly affecting the estimated effects of subsidized housing participation.

## **8 Concluding Comments**

In spite of the policy relevance of having a sound understanding of the effects of subsidized rental housing on long-term outcomes, the existing literature is clearly lacking. In this paper, we report results from a project that fills important gaps in this literature by estimating the long-term causal effects of public housing and voucher assisted housing participation as a teenager on adult earnings.

Our use of national data on housing assistance, households, and earnings from administrative records, censuses, and surveys at the U.S. Census Bureau makes these contributions possible. The data permit us to identify households with children between the ages of 13-18 in the year 2000, follow those children across a variety of settings of assisted and unassisted rental housing, and then to investigate their employment and earnings up to 10 years later.

We recognize that unobserved heterogeneity and the associated selection bias is an obstacle to estimating causal effects of housing. To overcome this issue, we exploit the very large sample size and longitudinal nature of the data and estimate household fixed effects models that identify the impact of assisted housing by exploiting variation within households. We also consider specifications that include time-varying household measures that



may vary across children, including parent's income and average neighborhood poverty, but these controls do not affect our estimates. One main finding is that the substantial negative effects of subsidized housing often found in the literature may be largely attributable to the negative selection of households entering assisted housing. A second main finding is that having controlled for unobservable heterogeneity with household fixed effects, subsidized housing participation as a teenager yields large positive effects on young adult earnings for females. For males we generally find no effect though in some cases the effect on earnings appears slightly negative.

The point estimates suggest that young adult females earn 14 percent more if they ever resided in voucher housing and 29 percent more if they ever resided in public housing. The corresponding estimates from the dose treatment indicate that each additional year of voucher-supported housing participation increases earnings by 6 percent for females while each additional year of public housing increases female earnings by 9 percent.

We disaggregate our sample by the race/ethnicity of households, to reflect the different contexts in which households select into assisted housing. We find results that differ considerably by race/ethnicity. IN particular, non-Hispanic Blacks, and especially non-Hispanic Black females, benefit more than Hispanics or non-Hispanic Whites. Thus, the main finding of a large, positive effect for females is primarily driven by the impact on Black females. The findings for Black teenagers are in strong contrast to the findings for non-Hispanic White and Hispanic teenagers. There were no significant positive or negative findings for Non-Hispanic White male or female earnings as a young adult. We also find few significant results for Hispanic teenagers.

We also investigated heterogeneous treatment effects by type of public housing (project size and project median income). We found no evidence that the effects of public housing on labor market outcomes varies along these dimensions. We also conducted robustness checks for measurement error and endogenous changes in family structure and found our results were robust to these concerns.

There remain a number of limitations of our analysis. First, our results apply to just two of the many subsidized housing programs, albeit the largest – public housing and housing vouchers. Second, our results might not be representative of all subsidized households (that is, households with younger children, and those with just one teenager). In this regard, our results pertain only teenagers between the ages 13 and 18. While this is a formative period, other research on siblings has shown that within-sibling differences in environment occurring early in childhood or even before birth can be critical. Future work should investigate whether exposure to subsidized housing during earlier periods of life has long-term implications as well.

## 9 Tables

**Table 1: Characteristics of Households Receiving  
Federal Rental Subsidies in the  
Form of Public Housing or Vouchers, 2000**

	<b>Public Housing</b>	<b>Voucher-Supported Housing</b>
Number of people per unit	2.3	2.7
Rent per month	\$202	\$226
Household income per year	\$10,000	\$10,600
Average months on waiting list	15	26
Average months since moved in	107	52
Percent of households where majority of income is derived from welfare	11%	12%
Percent of metropolitan area median income	25	23
Percent of households with children	45	61
Percent minority	69	61
Percent moved in past year	10	15
Percent with 0 or 1 bedrooms	48	25
Percent with 2 bedrooms	25	39
Percent with 3 or more bedrooms	27	35
<b>Total Households</b>	<b>1,282,099</b>	<b>1,817,360</b>

Source: HUDUSER, HUD Public Use Data.

**Table 2. Summary Statistics for Teenagers Aged 13-18 in 2000 in Renter Households Whose Parents Earned Less than Half of HUD-Specified Area Median Income in 2000**

	Individuals Aged 13-18 in 2000 With at least One Other Sibling Aged 13-18 in 2000				
	Total	In households not receiving any housing subsidy 1997-2005	Total	Teenagers who never lived in subsidized housing while aged 13-18	Teenagers who lived in subsidized housing while aged 13-18
	(1)	(2)	(3)	(4)	(5)
Household size in 2000	5.721	5.697	5.772	6.122	5.706
Age in 2000	15.381	15.424	15.287	15.537	15.24
Male	0.5	0.506	0.487	0.503	0.484
Black non-Hispanic household	0.312	0.232	0.482	0.463	0.486
Hispanic household	0.314	0.334	0.27	0.275	0.269
Other non-Hispanic household	0.08	0.086	0.068	0.065	0.068
White non-Hispanic household	0.294	0.348	0.18	0.196	0.177
Average block group percent poverty during teens	0.115	0.112	0.123	0.122	0.123
Average inverse hyperbolic sine of parents earnings during teens	8.006	8.219	7.546	7.662	7.524
Single-headed household	0.653	0.601	0.764	0.742	0.768
Ever public housing resident (13-18)	0.095	0	0.3	0	0.357
Ever a voucher recipient (13-18)	0.182	0	0.573	0	0.681
Years in public housing (13-18)	0.311	0	0.979	0	1.163
Years in voucher housing (13-18)	0.626	0	1.972	0	2.342
Total labor market earnings between 2008 and 2010	30000	33000	25000	25000	25000
Total number of quarters worked between 2008 and 2010	7.051	7.277	6.566	6.502	6.578
Proportion with any labor market earnings between 2008 and 2010	0.83	0.84	0.81	0.804	0.811
Number of observations (rounded)	520,000	358,000	162,000	25,000	137,000

SOURCE: Authors' tabulations of matched 2000 Census, HUD PIC, and LEHD files (see text).

NOTE: Excludes teenagers who lived in counties with at least one Public Housing Authority participating in HUD's Moving to Work Program.

**Table 3: The Effect of Teenage Residence in HUD-Subsidized Housing on Total 2008-2010 Earnings**

	Panel A: Dummy Treatment				Panel B: Dose Treatment				
	OLS	HFE	HFE with parents' earnings control	HFE with block group poverty control	OLS	HFE	HFE with parents' earnings control	HFE with block group poverty control	HFE with both earnings and poverty controls
Lives in a household receiving a housing voucher	-0.347***	0.135**	0.132**	0.135**	-0.077***	0.062***	0.061***	0.062***	0.061***
Male*Lives in a household receiving a housing voucher	0.021	0.046	0.046	0.046	0.006	0.012	0.012	0.012	0.012
Lives in public housing	-0.349***	-0.371***	-0.366***	-0.371***	-0.087***	-0.089***	-0.087***	-0.089***	-0.087***
Male*Lives in public housing	0.032	0.04	0.04	0.04	0.008	0.01	0.01	0.01	0.01
Male	-0.290***	0.292***	0.290***	0.289***	-0.069***	0.088***	0.087***	0.087***	0.086***
Natural log of average parents' earnings (ages 13-18)	0.027	0.059	0.059	0.059	0.007	0.018	0.018	0.018	0.018
Male*Natural log of average parents' earnings (ages 13-18)	-0.336***	-0.360***	-0.354***	-0.354***	-0.079***	-0.086***	-0.084***	-0.084***	-0.083***
Average block group poverty (ages 13-18)	0.042	0.053	0.053	0.053	0.012	0.014	0.014	0.014	0.014
Constant	-0.461***	-0.407***	-0.485***	-0.374***	-0.472***	-0.422***	-0.502***	-0.387***	-0.469***
R-squared	0.029	0.036	0.05	0.044	0.029	0.036	0.05	0.044	0.058
			0.013	0.013			0.013	0.013	0.013
			0.016	0.016			0.016	0.016	0.016
			0.009*	0.009*			0.010*	0.009*	0.009*
			0.004	0.004			0.004	0.004	0.004
			-0.127	-0.149			-0.123	-0.145	-0.145
			0.468	0.468			0.468	0.468	0.468
			-0.293	-0.249			-0.313	-0.268	-0.268
			0.241	0.241			0.241	0.241	0.241
	8.900***	8.722***	8.619***	8.737***	8.881***	8.704***	8.601***	8.719***	8.617***
	0.02	0.025	0.134	0.059	0.02	0.025	0.134	0.059	0.145
	0.009	0.006	0.006	0.006	0.008	0.006	0.006	0.006	0.006

NOTES: OLS = Ordinary Least Squares. HFE = Household Fixed Effects. Number of observations = 520,000, rounded to the nearest thousand. See text for description of the sample. The dependent variable in each column is the inverse hyperbolic sine of total earnings between 2008 and 2010. All columns include controls for age, sex, age by sex, and household race by sex; race and ethnicity is determined by the reference person of the household as reported on the 2000 Census. Column 3 also includes a control for the inverse hyperbolic sine of average parents' annual earnings while the teenager was 13-18 years of age. Column 4 also includes a control for the average block group percent poverty experienced while a teenager. Column 5 includes both controls. Standard Errors below the estimates. Based on the authors' tabulations from matched Census 2000-LEHD-PIC file. \*\*\* p<=0.01, \*\* p<=0.05, \* p<=0.10.

**Table 4: The Effect of Teenage Residence in HUD-Subsidized Housing on Total 2008-2010 Earnings  
Non-Hispanic White Households Only**

	Panel A: Dummy Treatment				Panel B: Dose Treatment					
	OLS	HFE	HFE with parents' earnings control	HFE with block group poverty control	HFE with both earnings and poverty controls	OLS	HFE	HFE with parents' earnings control	HFE with block group poverty control	HFE with both earnings and poverty controls
Lives in a household receiving a housing voucher	-0.435***	0.005	0.002	0.009	0.006	-0.117***	0.033	0.032	0.034	0.033
Male*Lives in a household receiving a housing voucher	0.042	0.088	0.088	0.087	0.088	0.012	0.026	0.026	0.026	0.026
Lives in public housing	-0.312***	0.076	0.076	0.076	0.076	-0.016	0.021	-0.006	-0.01	-0.009
Male*Lives in public housing	0.067	0.125	0.126	0.126	0.126	0.017	0.016	0.021	0.021	0.021
Natural log of average parents' earnings (ages 13-18)	0.298**	0.246*	0.250*	0.235	0.238	-0.081***	0.045	0.045	0.046	0.046
Male*Natural log of average parents' earnings (ages 13-18)	0.098	0.122	0.122	0.122	0.122	0.021	0.05	0.051	0.047	0.048
Average block group poverty (ages 13-18)	-0.224***	-0.185***	-0.256**	-0.266***	-0.338***	0.029	0.035	0.036	0.036	0.036
Male*Average block group poverty (ages 13-18)	0.045	0.055	0.079	0.074	0.094	-0.218***	-0.177**	-0.247**	-0.262***	-0.332***
Constant	8.813***	8.732***	8.773***	8.760***	8.800***	0.045	0.055	0.079	0.074	0.094
R-squared	0.031	0.037	0.219	0.097	0.237	-0.005	0.027	-0.005	-0.005	-0.005
	0.003	0.003	0.003	0.003	0.003	0.027	0.009	0.027	0.027	0.027
						0.009	0.007	0.009	0.009	0.009
						0.007	0.007	0.007	0.007	0.007
						-0.294	-0.312	-0.306	-0.312	-0.306
						0.94	0.94	0.94	0.94	0.94
						0.87	0.868	0.868	0.903	0.901
						0.561	0.561	0.561	0.56	0.56
						8.713***	8.713***	8.755***	8.742***	8.783***
						0.031	0.038	0.219	0.097	0.237
						0.003	0.003	0.003	0.003	0.003

NOTES: See Table 3. Number of observations = 197,000, rounded to the nearest thousand

**Table 5: The Effect of Teenage Residence in HUD-Subsidized Housing on Total 2008-2010 Earnings  
Non-Hispanic Black Households Only**

	Panel A: Dummy Treatment				Panel B: Dose Treatment					
	OLS	HFE	HFE with parents' earnings control	HFE with block group poverty control	HFE with both earnings and poverty controls	OLS	HFE	HFE with parents' earnings control	HFE with block group poverty control	HFE with both earnings and poverty controls
Lives in a household receiving a housing voucher	-0.124***	0.145*	0.146*	0.140*	0.142*	-0.017*	0.061***	0.061***	0.060***	0.060***
Male*Lives in a household receiving a housing voucher	0.03	0.066	0.066	0.066	0.066	0.007	0.018	0.018	0.018	0.018
Lives in public housing	-0.174***	-0.229***	-0.232***	-0.218***	-0.223***	-0.030*	-0.033*	-0.033*	-0.030*	-0.031*
Male*Lives in public housing	0.049	0.061	0.061	0.061	0.061	0.012	0.015	0.015	0.015	0.015
Natural log of average parents' earnings (ages 13-18)	-0.176***	0.177*	0.178*	0.180*	0.181*	-0.035***	0.058*	0.059*	0.060*	0.061*
Male*Natural log of average parents' earnings (ages 13-18)	0.035	0.078	0.078	0.078	0.078	0.009	0.024	0.024	0.024	0.024
Average block group poverty (ages 13-18)	-0.048	-0.05	-0.058	-0.06	-0.065	0.007	0.012	0.009	0.008	0.007
Male*Average block group poverty (ages 13-18)	0.057	0.073	0.073	0.073	0.073	0.015	0.019	0.019	0.02	0.02
Constant	-1.356***	-1.253***	-1.101***	-1.361***	-1.206***	-1.385***	-1.300***	-1.151***	-1.410***	-1.259***
R-squared	0.052	0.065	0.088	0.082	0.107	0.052	0.064	0.088	0.081	0.106
			0.066*	0.065*	0.065*			0.066*		0.065*
			0.031	0.031	0.031			0.031		0.031
			-0.019*	-0.017*	-0.017*			-0.019*		-0.016*
			0.007	0.007	0.007			0.007		0.007
				-0.706	-0.632				-0.729	-0.658
				0.72	0.72				0.72	0.72
				0.861*	0.706				0.881*	0.731
				0.408	0.413				0.408	0.413
	8.821***	8.618***	8.099***	8.707***	8.185***	8.792***	8.592***	8.073***	8.683***	8.162***
	0.033	0.044	0.243	0.1	0.259	0.033	0.045	0.243	0.1	0.259
	0.029	0.033	0.033	0.033	0.033	0.029	0.033	0.033	0.033	0.033

NOTES: See Table 3. Number of observations = 162,000, rounded to the nearest thousand

**Table 6: The Effect of Teenage Residence in HUD-Subsidized Housing on Total 2008-2010 Earnings  
Hispanic Households Only**

	Panel A: Dummy Treatment					Panel B: Dose Treatment				
	OLS	HFE	HFE with parents' earnings control	HFE with block group poverty control	HFE with both earnings and poverty controls	OLS	HFE	HFE with parents' earnings control	HFE with block group poverty control	HFE with both earnings and poverty controls
Lives in a household receiving a housing voucher	-0.840***	-0.159	-0.174	-0.159	-0.174	-0.194***	0.005	0.001	0.005	0.001
Male*Lives in a household receiving a housing voucher	0.049	0.101	0.101	0.101	0.101	0.013	0.026	0.026	0.026	0.026
Lives in public housing	-0.041	-0.025	0.006	-0.025	0.006	-0.021	-0.019	-0.011	-0.019	-0.011
Male*Lives in public housing	0.071	0.082	0.083	0.082	0.083	0.018	0.021	0.021	0.021	0.022
Natural log of average parents' earnings (ages 13-18)	-0.693***	0.108	0.096	0.107	0.097	-0.167***	0.033	0.03	0.033	0.03
Parents' earnings (ages 13-18)	0.063	0.133	0.133	0.133	0.133	0.017	0.038	0.038	0.038	0.038
Average block group poverty (ages 13-18)	0.027	-0.018	0.007	-0.017	0.004	0	-0.017	-0.011	-0.017	-0.012
Male*Average block group poverty (ages 13-18)	0.091	0.11	0.11	0.11	0.11	0.024	0.03	0.03	0.03	0.03
Constant	-0.025	0.001	-0.241*	0.007	-0.251*	-0.011	0.013	-0.222*	0.016	-0.235*
R-squared	0.058	0.072	0.103	0.087	0.116	0.058	0.071	0.102	0.087	0.115
			-0.003		-0.003			-0.002		-0.002
			0.032		0.032			0.032		0.032
			0.028***		0.028***			0.027**		0.027***
			0.008		0.008			0.008		0.008
			0.031		-0.038			0.01		-0.057
			0.84		0.841			0.84		0.841
			-0.046		0.071			-0.021		0.094
			0.396		0.396			0.396		0.397
	9.118***	8.905***	8.933***	8.901***	8.939***	9.086***	8.875***	8.897***	8.873***	8.905***
	0.04	0.049	0.266	0.118	0.288	0.04	0.05	0.266	0.118	0.288
	0.009	0.002	0.003	0.002	0.003	0.008	0.002	0.003	0.002	0.003

NOTES: See Table 3. Number of observations = 123,000, rounded to the nearest thousand



**Table 7: Summary of the Household Fixed Effects Estimates  
By Gender and Race/Ethnicity**

	Dummy			Dose		
	Housing Voucher (HV) Treatment Effect	Public Housing (PH) Treatment Effect	Are Subsidy Effects Different? (HV vs. PH)	Housing Voucher (HV) Treatment Effect	Public Housing (PH) Treatment Effect	Are Subsidy Effects Different? (HV vs. PH)
<b><u>All Households</u></b>						
Females (F)	0.135**	0.292***	Yes**	0.062***	0.088***	No
	0.046	0.059		0.012	0.018	
Males (M)	-0.236***	-0.068	Yes*	-0.027**	0.002	No
	0.048	0.064		0.013	0.019	
Are Subsidy Effects Different? (F vs. M)	Yes***	Yes***		Yes***	Yes***	
<b><u>Non-Hispanic White Households</u></b>						
Females	0.005	-0.082	No	0.033	0.016	No
	0.088	0.125		0.026	0.045	
Males	0.003	0.165	No	0.025	0.066	No
	0.09	0.136		0.026	0.048	
Are Subsidy Effects Different? (F vs. M)	No	Yes*		No	No	
<b><u>Non-Hispanic Black Households</u></b>						
Females	0.145*	0.177*	No	0.061***	0.058*	Yes*
	0.066	0.078		0.018	0.024	
Males	-0.084	0.126	Yes**	0.028	0.070***	No
	0.071	0.087		0.019	0.025	
Are Subsidy Effects Different? (F vs. M)	Yes***	No		No	No	
<b><u>Hispanic Households</u></b>						
Females	-0.159	0.108	No	0.005	0.033	No
	0.101	0.133		0.026	0.038	
Males	-0.184*	0.089	Yes*	-0.014	0.016	No
	0.104	0.137		0.026	0.039	
Are Subsidy Effects Different? (F vs. M)	No	No		No	No	

NOTES: All columns present household fixed effects estimates of the effect of subsidized housing participation as a teenager on the inverse hyperbolic sine of total earnings 2008-2010. Estimates do not control for parents' earnings as a teenager or average block group percent poverty as a teenager but include a male indicator and a full set of age in years by male fixed effects. Observations rounded to the nearest thousand (All: 521,000; non-Hispanic White: 197,000; non-Hispanic Black: 162,000; Hispanic: 123,000). Standard errors are below the estimates. \*\*\* p<=0.01, \*\* p<=0.05, \* p<=0.10.

**Table 8: Subsidized Housing Residence and Adult Earnings by Household Race and Gender  
Allowing for a Differential Effect of Large Public Housing**

	<i>Black Households</i>		<i>Hispanic Households</i>		<i>White Households</i>		<i>All Households</i>	
	Dummy	Dose	Dummy	Dose	Dummy	Dose	Dummy	Dose
Lives in a household receiving a housing voucher	0.145*	0.061***	-0.157	0.005	0.005	0.033	0.136**	0.062***
Male*Lives in a household receiving a housing voucher	0.066	0.018	0.101	0.026	0.088	0.026	0.046	0.012
Lives in public housing	-0.229***	-0.033*	-0.028	-0.019	-0.002	-0.008	-0.372***	-0.089***
Male*Lives in public housing	0.061	0.015	0.082	0.021	0.076	0.021	0.04	0.01
Lives in a large public housing project	0.172	0.076**	-0.005	0.019	-0.111	0.01	0.233***	0.093***
Male*Lives in a large public housing project	0.088	0.028	0.154	0.044	0.13	0.047	0.066	0.021
Lives in a large public housing project	-0.024	0.004	0.158	0.015	0.257*	0.063	-0.266***	-0.077***
Male*Lives in a large public housing project	0.086	0.023	0.138	0.038	0.126	0.037	0.062	0.017
Lives in a large public housing project	0.011	-0.057	0.28	0.032	0.289	0.071	0.192	-0.016
Male*Lives in a large public housing project	0.147	0.051	0.257	0.079	0.425	0.166	0.12	0.041
Lives in a large public housing project	-0.077	0.023	-0.432*	-0.076	-0.093	-0.156	-0.285**	-0.028
Male*Lives in a large public housing project	0.141	0.043	0.213	0.063	0.422	0.131	0.109	0.033
Constant	8.618***	8.591***	8.905***	8.875***	8.732***	8.713***	8.722***	8.704***
Observations	162,000	162,000	123,000	123,000	197,000	197,000	520,000	520,000

NOTES: Authors' tabulations from matched Census 2000-LEHD-PIC file. Coefficients from Household Fixed Effects regressions with the inverse hyperbolic sine of earnings between 2008 and 2010 as the dependent variable. In addition to the controls discussed in Tables 3-7, Table 8 allows for two different types of public housing defined by project total population. Large public housing projects represent the top quartile of projects with respect to population between 1997 and 2005.

**Table 9: Subsidized Housing Residence and Adult Earnings by Household Race and Gender  
Allowing for a Differential Effect of Low Income Public Housing**

	Black Households		Hispanic Households		White Households		All Households	
	Dummy	Dose	Dummy	Dose	Dummy	Dose	Dummy	Dose
Lives in a household receiving a housing voucher	0.144*	0.061***	-0.158	0.005	0.005	0.033	0.133**	0.062***
Male*Lives in a household receiving a housing voucher	0.066	0.018	0.101	0.026	0.088	0.026	0.046	0.012
Lives in public housing	-0.227***	-0.032*	-0.026	-0.019	-0.002	-0.008	-0.368***	-0.088***
Male*Lives in public housing	0.061	0.015	0.082	0.021	0.076	0.021	0.04	0.01
Lives in a low-income public housing project	0.112	0.033	0.087	0.035	-0.144	0.006	0.176*	0.057**
Male*Lives in a low-income public housing project	0.093	0.028	0.144	0.041	0.142	0.05	0.068	0.021
Lives in a household receiving a housing voucher	0.02	0.026	-0.032	-0.024	0.242	0.051	-0.229***	-0.058***
Male*Lives in public housing	0.09	0.023	0.118	0.031	0.14	0.039	0.062	0.016
Lives in a low-income public housing project	0.153	0.088	0.099	-0.029	0.213	0.056	0.348***	0.144**
Male*Lives in a low-income public housing project	0.125	0.051	0.283	0.124	0.241	0.134	0.103	0.044
Lives in a low-income public housing project	-0.174	-0.054	0.097	0.079	0.032	-0.004	-0.438***	-0.143***
Constant	0.134	0.047	0.293	0.114	0.278	0.124	0.109	0.04
	8.619***	8.592***	8.906***	8.875***	8.732***	8.713***	8.724***	8.705***
Observations	162,000	162,000	123,000	123,000	197,000	197,000	520,000	520,000

NOTES: Coefficients from Household Fixed Effects regressions with the inverse hyperbolic sine of earnings between 2008 and 2010 as the dependent variable. In addition to the controls discussed in Tables 3-7, Table 9 allows for two different types of public housing defined by median household income per project. Low-income public housing projects represent the bottom quartile of projects with respect to person-weighted median household income in each year between 1997 and 2005.

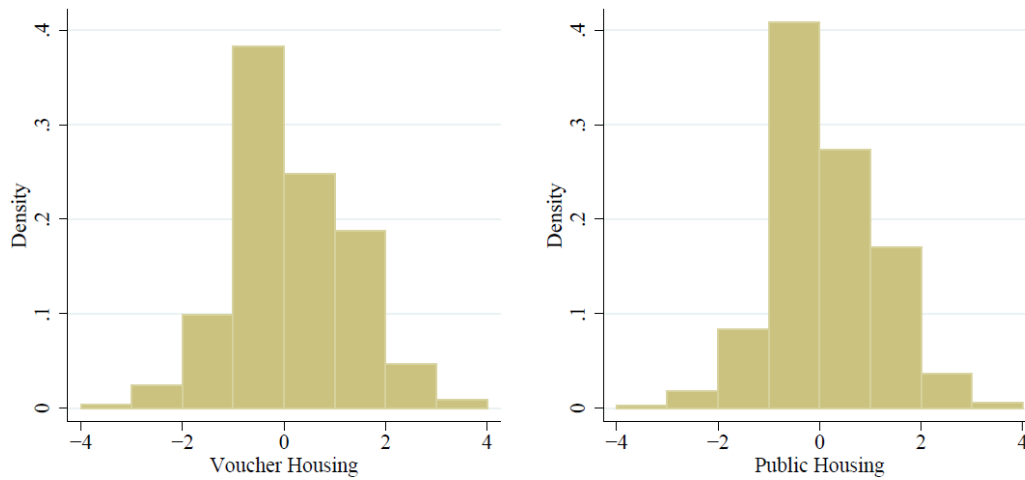
**Table 10: Subsidized Housing Residence and Adult Earnings  
Using Actual Participation, Predicted Participation, and Actual Instrumented by Predicted Participation**

	Dummy			Dose		
	HFE	HFE Predicted Treatment	HFE IV (Predicted for Actual Treatment)	HFE	HFE Predicted Treatment	HFE IV (Predicted for Actual Treatment)
Lives in a household receiving a housing voucher	0.135** 0.046	0.168* 0.066	0.194* 0.099	0.062*** 0.012	0.045** 0.017	0.054* 0.021
Male*Lives in a household receiving a housing voucher	-0.371*** 0.04	-0.403*** 0.04	-0.474*** 0.047	-0.089*** 0.01	-0.090*** 0.01	-0.108*** 0.012
Lives in public housing	0.292*** 0.059	0.217* 0.091	0.278 0.15	0.088*** 0.018	0.043 0.025	0.052 0.033
Male*Lives in public housing	-0.360*** 0.053	-0.383*** 0.054	-0.460*** 0.065	-0.086*** 0.014	-0.088*** 0.014	-0.105*** 0.017

NOTES: Number of observations = 520,000. Table 10 presents only the coefficients on the two housing subsidy measures and their interactions with a male indicator from six different specifications. The HFE columns repeats the estimates from our main household fixed effects specifications to simplify comparison. The HFE Predicted columns present household fixed effects regressions where we define participation in subsidized housing using the 2000 head of household's movements in and out of subsidized housing, as well as an individual's age, to define program participation, instead of observed participation from the teenager's administrative record. The HFE IV columns use the predicted treatment as an instrument for the teenager's actual treatment in a fixed effects instrumental variables specification. In both cases, the first stage F-statistics are well above conventional thresholds for weak instruments. In all columns the inverse hyperbolic sine of earnings between 2008 and 2010 is the dependent variable.


## 10 Figures

Figure 1: Distribution of Within-Household Differences in Subsidized Housing Participation



Note: Figure displays the distribution of within-household differences in public housing and housing voucher participation for teenagers in the main sample. Within-household differences are topcoded to have an absolute value no greater than four and individuals from households with no differences in program participation are omitted. Of individuals in households with some voucher housing participation, 0.417 have no within-household variation. Of individuals in households with some public housing participation, 0.693 have no within-household variation. Each bin represents a one year difference in program participation.

Figure 2: Days Spent on Waitlist for Program Admission in 2000



Note: Figure displays the distribution of days spent on the waiting list before admission for households found in both public and voucher housing in the year 2000. The sample is limited to households with non-missing admission and waitlist information who gained admission to their program no earlier than 1995. 0.116 of public housing households spent >1 year and 0.033 spent >2 years on a waitlist prior to admission. 0.287 of voucher housing households spent >1 year and 0.108 spent >2 years on a waitlist prior to admission.

# 11 Appendix A: Major U.S. Subsidized Rental Housing Programs

There are a wide variety of subsidized housing programs in the United States. Beginning in the 1930s, the U.S. government built public housing projects, and for decades, the program continued to be the primary means of federal assistance for rental housing. The Housing Act of 1949 introduced income limits and “Fair Market Rents” along with subsidies that would incentivize private development of low-cost housing and were further expanded in the late 1960s. In the 1980s, production was drastically reduced as housing assistance became a more decentralized effort, and no federal public housing has been built since 1981. A “regime change” in the mid-1980s additionally introduced even stricter requirements to focus assistance on the poorest households. There were about 1.4 million public housing units in 1990, falling to just under 1.3 million in 2000, and about 1.1 million in 2008. The reduction in these numbers reflects demolition of the worst-performing projects starting in the 1990s. In these cases, under the HOPE VI program, tenants are typically given housing vouchers to find housing elsewhere (Popkin et al. 2004). Today, over 3,000 Public Housing Authorities administer public housing projects, mostly for the very poor and typically neighborhoods that are predominantly low-income.

The Housing Choice Voucher Program (HCVP) provides direct rental assistance to housing tenants through vouchers. The Section 8 New Construction and Substantial Rehabilitation project-based subsidy program assists owners of housing units so that they may charge affordable rents; it accounted for almost 900,000 units in 2000. Note that these households are much smaller and live in smaller dwellings than their counterparts in residing public housing or receiving vouchers. This reflects in part the large share of elderly occupants.

While Section 8 subsidized housing began as project-based housing subsidy in 1974 and at that time was based on new construction, now much of the housing historically referred

to as Section 8 housing is found in the tenant-based HCVP program. HCVP has developed more recently and is solely a demand-side, tenant-based subsidy program. Stemming from the ambitious Experimental Housing Allowance Program of the 1970s (see Friedman and Weinberg, 1982, 1983) this program brings a different perspective to housing policy by separating itself from new production. Rather than choosing among specific subsidized housing locations, voucher recipients may live in any structurally adequate rental housing in a specified rent and size range, with the Federal subsidy making the unit affordable. Public Housing Authorities may to allocate up to 20 percent of their HCVP funds for project-based vouchers that are tied to specific private housing developments, rather than to the tenant. Tenant vouchers can be used by those wishing to live in Low Income Housing Tax Credit housing (described below) and thus there is the potential for multiple types of subsidies for a given unit. This program provides anonymity and a choice of locations, although landlord willingness to participate limits its extent. There were about 1.1 million voucher households in 1990, growing dramatically to 1.8 million in 2000, and continuing to grow. Currently, over 30 percent of U.S. subsidized housing is provided by vouchers.

The Low Income Housing Tax Credit (LIHTC) program began with the 1986 Tax Reform Act, and was expanded by 40 percent in 2001. Unlike the “deep subsidies” provided by the other three programs discussed here, LIHTC provides “shallow subsidies” in that no ongoing operating costs are covered by the government. In this program, the U.S. government (through the Internal Revenue Service), provides tax credits to for-profit and non-profit developers to build income-restricted housing. In 1990, there were about 140,000 units, growing to almost 1 million in 2000, and growing further to almost 1.7 million units in 2008. While LIHTC housing has significant income limits for eligibility, this program does not provide housing for the very poor. Another concern raised about the LIHTC program is that it may crowd out nearby private investment in affordable rental housing, as Eriksen and Rosenthal [2010] find.



## References

- Country water action: Indonesia. simple solution for drinking water makes big difference, March 2006. URL <http://www.adb.org/water/actions/ino/simple-solution.asp>.
- Daniel Aaronson. Using sibling data to estimate the impact of neighborhoods on children's educational outcomes. *The Journal of Human Resources*, 33(2):915–946, 1998. URL <http://www.jstor.org/stable/146403>.
- John Abowd, John Haltiwanger, and Julia Lane. Integrated longitudinal employee-employer data for the united states. *American Economic Review*, 94(2):224–229, 2004.
- Kathleen Abu-Saad and Drora Fraser. Maternal nutrition and birth outcomes. *Epidemiologic Reviews*, 32:5–25, 2010. doi: 10.1093/epirev/mxq001.
- A. R. A. Adegboye and B. L. Heitmann. Accuracy and correlates of maternal recall of birthweight and gestational age. *BJOG: An International Journal of Obstetrics & Gynaecology*, 115(7):886–893, June 2008. ISSN 14700328. doi: 10.1111/j.1471-0528.2008.01717.x. URL <http://dx.doi.org/10.1111/j.1471-0528.2008.01717.x>.
- Achyuta R. Adhvaryu and Anant Nyshadham. Endowments at birth and parents' investments in children. Technical report, Yale University Working Paper, 2012.
- Naeem Ahmed, Matthew Brzozowski, and Thomas F. Crossley. Measurement Errors in Recall Food Expenditure Data. Quantitative Studies in Economics and Population Research Reports, McMaster University 396, McMaster University, October 2005. URL <http://ideas.repec.org/p/mcm/qsepr/396.html>.
- Anna Aizer and Flavio Cunha. Child endowments, parental investments and the development of human capital: Evidence from siblings. Technical report, Brown University Working Paper, 2012.

- Richard Akresh, Emilie Bagby, Damien de Walque, and Harounan Kazianga. Child labor, schooling, and child ability. *Economic Development and Cultural Change*, 61(1):157–186, 2012.
- Douglas Almond and Janet Currie. Human capital development before age five. Technical report, NBER, 2010.
- Douglas Almond and Janet Currie. Killing me softly: The fetal origins hypothesis. *Journal of Economic Perspectives*, 25(3):153–72, Summer 2011. URL <http://ideas.repec.org/a/aea/jecper/v25y2011i3p153-72.html>.
- Douglas Almond and Bhashkar Mazumder. The 1918 influenza pandemic and subsequent health outcomes: An analysis of sipp data. *American Economic Review*, 95(2):258–262, May 2005. URL <http://ideas.repec.org/a/aea/aecrev/v95y2005i2p258-262.html>.
- Douglas Almond and Bhashkar Mazumder. Health capital and the prenatal environment: The effect of ramadan observance during pregnancy. *American Economic Journal: Applied Economics*, 3(4):56–85, October 2011. URL <http://ideas.repec.org/a/aea/aejapp/v3y2011i4p56-85.html>.
- Douglas Almond, Kenneth Y. Chay, and David S. Lee. The costs of low birth weight. *The Quarterly Journal of Economics*, 120(3):1031–1083, August 2005. URL <http://ideas.repec.org/a/tpr/qjecon/v120y2005i3p1031-1083.html>.
- Douglas Almond, Lena Edlund, and Mårten Palme. Chernobyl’s subclinical legacy: Prenatal exposure to radioactive fallout and school outcomes in sweden. *The Quarterly Journal of Economics*, 124(4):1729–1772, November 2009. URL <http://ideas.repec.org/a/tpr/qjecon/v124y2009i4p1729-1772.html>.
- Douglas Almond, Hilary W. Hoynes, and Diane Whitmore Schanzenbach. Inside the war on poverty: The impact of food stamps on birth outcomes. *The Review of Economics*

- and Statistics*, 93(2):387–403, December 2011. URL <http://www.nber.org/papers/w14306>.
- Joshua D. Angrist and William N. Evans. Children and their parents' labor supply: Evidence from exogenous variation in family size. *The American Economic Review*, 88(3): pp. 450–477, 1998. ISSN 00028282. URL <http://www.jstor.org/stable/116844>.
- Joshua D. Angrist and Alan B. Krueger. The effect of age at school entry on educational attainment: An application of instrumental variables with moments from two samples. *Journal of the American Statistical Association*, 87(418):328–336, June 1992. URL <http://www.jstor.org/stable/2290263>.
- Fred Arnold, Minja Kim Choe, and T. K. Roy. Son preference, the family-building process and child mortality in india. *Population Studies*, 52(3):pp. 301–315, 1998. ISSN 00324728. URL <http://www.jstor.org/stable/2584732>.
- Michael Baker and Kevin Milligan. Boy-girl differences in parental time investments: Evidence from three countries. NBER Working Papers 18893, National Bureau of Economic Research, Inc, March 2013. URL <http://ideas.repec.org/p/nbr/nberwo/18893.html>.
- Neil Bania, Claudia Coulton, and Laura Leete. Public housing assistance, public transportation, and the welfare-to-work transition. *Cityscape: A Journal of Policy Development and Research*, 6(2):7–44, 2003.
- Silvia H. Barcellos, Leandro Carvalho, and Adriana Lleras-Muney. Child gender and parental investments in india: Are boys and girls treated differently? Working Paper 17781, National Bureau of Economic Research, January 2012. URL <http://www.nber.org/papers/w17781>.
- David J.P. Barker and Phillipa M. Clark. Fetal undernutrition and disease in later life. *Reviews of Reproduction*, 2:105–112, 1997.

- DJ Barker, PD Winter, C Osmond, B Margetts, and SJ Simmonds. Weight in infancy and death from ischaemic heart disease. *Lancet*, 2(8663):577–580, 1989.
- D.J.P. Barker. *Mothers, Babies, and Health in Later Life*. Churchill Livingstone, 1998.
- Gary S. Becker. An economic analysis of fertility. In *Demographic and Economic Change in Developed Countries*, NBER Chapters, pages 209–240. National Bureau of Economic Research, Inc, December 1960. URL <http://ideas.repec.org/h/nbr/nberch/2387.html>.
- Gary S Becker and H Gregg Lewis. On the interaction between the quantity and quality of children. *Journal of Political Economy*, 81(2):S279–88, Part II, 1973. URL <http://ideas.repec.org/a/ucp/jpolec/v81y1973i2ps279-88.html>.
- Gary S Becker and H. Gregg Lewis. Interaction between quantity and quality of children. In Theodore W. Schultz, editor, *Economics of the Family: Marriage, Children, and Human Capital*. UMI, 1974.
- Gary S Becker and Nigel Tomes. Child endowments and the quantity and quality of children. *Journal of Political Economy*, 84(4):S143–62, August 1976. URL <http://ideas.repec.org/a/ucp/jpolec/v84y1976i4ps143-62.html>.
- Jere R Behrman, Robert A Pollak, and Paul Taubman. Parental preferences and provision for progeny. *Journal of Political Economy*, 90(1):52–73, February 1982. URL <http://ideas.repec.org/a/ucp/jpolec/v90y1982i1p52-73.html>.
- Yoram Ben-Porath and Finis Welch. Do sex preferences really matter? *The Quarterly Journal of Economics*, 90(2):285–307, May 1976. URL <http://ideas.repec.org/a/tpr/qjecon/v90y1976i2p285-307.html>.
- Sonia Bhalotra and Arthur van Soest. Birth-spacing, fertility and neonatal mortality in india: Dynamics, frailty, and fecundity. *Journal of Econometrics*, 143(2):274–290, April

2008. URL <http://ideas.repec.org/a/eee/econom/v143y2008i2p274-290.html>.
- Prashant Bharadwaj and Leah K. Lakdawala. Discrimination begins in the womb: Evidence of sex-selective prenatal investments. *Journal of Human Resources*, 48(1):71–113, 2013. URL <http://ideas.repec.org/a/uwp/jhriss/v48y2013i1p71-113.html>.
- Prashant Bharadwaj, Katrine Vellesen Loken, and Christopher Neilson. Early life health interventions and academic achievement. Technical report, Working Paper, 2012.
- P. N. Mari Bhat and A. J. Francis Xavier. Fertility decline and gender bias in northern india. *Demography*, 40(4):pp. 637–657, 2003. ISSN 00703370. URL <http://www.jstor.org/stable/1515201>.
- Marianne P. Bitler and Janet Currie. Does wic work? the effects of wic on pregnancy and birth outcomes. *Journal of Policy Analysis and Management*, 24(1):pp. 73–91, 2005a. ISSN 02768739. URL <http://www.jstor.org/stable/3326170>.
- Marianne P. Bitler and Janet Currie. The changing association between prenatal participation in wic and birth outcomes in new york city: What does it mean? *Journal of Policy Analysis and Management*, 24(4):pp. 687–690, 2005b. ISSN 02768739. URL <http://www.jstor.org/stable/30162674>.
- Sandra E. Black, Paul J. Devereux, and Kjell G. Salvanes. From the cradle to the labor market? the effect of birth weight on adult outcomes. *The Quarterly Journal of Economics*, 122(1):pp. 409–439, 2007. ISSN 00335533. URL <http://www.jstor.org/stable/25098846>.
- Sandra E. Black, Paul J. Devereux, and Kjell G. Salvanes. Small family, smart family? family size and the iq scores of young men. *The Journal of Human Resources*, 45: 33–58, 2010. URL <http://jhr.uwpress.org/content/45/1/33.refs>.

- Hoyt Bleakley. Malaria eradication in the americas: A retrospective analysis of childhood exposure. *American Economic Journal: Applied Economics*, 2:1–45, 2010. doi: 10.1257/app.2.2.1.
- Peter Boomgaard. Bridewealth and birth control: Low fertility in the indonesian archipelago, 1500-1900. *Population and Development Review*, 29(2):pp. 197–214, 2003. ISSN 00987921. URL <http://www.jstor.org/stable/3115225>.
- Jennifer Brown. Rural women’s land rights in java, indonesia: Strengthened by family law, but weakweak by land regulation. *Pacific Rim Law & Policy Journal*, 12, 2003.
- Martin Browning, Thomas F. Crossley, and Joachim Winter. The measurement of household consumption expenditures. *Annual Review of Economics*, 6(1):475–501, 2014. doi: 10.1146/annurev-economics-080213-041247. URL <http://dx.doi.org/10.1146/annurev-economics-080213-041247>.
- Kasey Buckles and Daniel Hungerman. Season of birth and later outcomes: Old questions, new answers. Technical report, NBER Working Paper, 2010.
- Kasey S. Buckles and Elizabeth L. Munnich. Birth spacing and sibling outcomes. *Journal of Human Resources*, 47(3):613–642, 2012. URL <http://ideas.repec.org/a/uwp/jhriss/v46y2012iii1p613-642.html>.
- John B Burbidge, Lonnie Magee, and A Leslie Robb. Alternative transformations to handle extreme values of the dependent variable. *Journal of the American Statistical Association*, 83(401):123–127, 1988.
- Julio Caceres-Delpiano. The impacts of family size on investment in child quality. *Journal of Human Resources*, 41:738–754, 2006.
- Deven Carlson, Robert Haveman, Thomas Kaplan, and Bar. Long-term effect of pub-

- lic low-income housing voucher on neighborhood quality and household composition. *Journal of Housing Economics*, 21:101–120, 2012a.
- Deven Carlson, Robert Haveman, Thomas Kaplan, and Barbara Wolfe. Long term earnings and employment effects of housing voucher receipt. *Journal of Urban Economics*, 71(1):128–150, 2012b.
- Eliana Carranza. Islamic inheritance law, son preference and fertility behavior of Muslim couples in Indonesia. Policy Research Working Paper Series 5972, The World Bank, February 2012. URL <http://ideas.repec.org/p/wbk/wbrwps/5972.html>.
- Anne Case, Angela Fertig, and Christina Paxson. The lasting impact of childhood health and circumstance. *Journal of Health Economics*, 24(2):365–389, March 2005. URL <http://ideas.repec.org/a/eee/jhecon/v24y2005i2p365-389.html>.
- Laura E Caulfield, Mercedes de Onis, Monika Blossner, and Rober E Black. Undernutrition as an underlying cause of child deaths associated with diarrhea, pneumonia, malaria, and measles. *American Journal of Clinical Nutrition*, 80(1):193–198, 2004.
- S.M. Ceesay, A.M. Prentice, T.J. Cole, F. Ford, L.T. Weaver, E.M. Poskitt, and R.G. Whitehead. Effects on birth weight and perinatal mortality of maternal dietary supplements in rural gambia: 5 year randomized controlled trial. *British Medical Journal*, 315:786–790, 1997. URL <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2127544/>.
- Raj Chetty, John N. Friedman, Nathaniel Hilger, Emmanuel Saez, Diane Whitmore Schanzenbach, and Danny Yagan. How does your kindergarten classroom affect your earnings? evidence from project star. *The Quarterly Journal of Economics*, 126(4):1593–1660, 2011. URL <http://ideas.repec.org/a/oup/qjecon/v126y2011i4p1593-1660.html>.
- Raj Chetty, John N Friedman, and Emmanuel Saez. Using differences in knowledge across

- neighborhoods to uncover the impacts of the eite on earnings. *American Economic Review*, 103(7):2683–2721, 2013.
- Raj Chetty, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez. Where is the land of opportunity? the geography of intergenerational mobility in the united states. *Quarterly Journal of Economics*, 129(4):1553–1623, 2014.
- CIA, 2012. URL <https://www.cia.gov/library/publications/the-world-factbook/geos/id.html>.
- CIA. The world factbook: Indonesia. website, June 2014. URL <https://www.cia.gov/library/publications/the-world-factbook/geos/id.html>.
- Shelley Clark. Son preference and sex composition of children: Evidence from india. *Demography*, 37(1):95–108, February 2000. URL <http://ideas.repec.org/a/spr/demogr/v37y2000i1p95-108.html>.
- TF Crossley and JK Winter. Asking households about expenditures: What have we learned? In C Carroll, TF Crossley, and J Sablehaus, editors, *Improving the measurement of consumer expenditures*. University of Chicago Press, 2015.
- Guillermo Cruces and Sebastian Galiani. Fertility and female labor supply in Latin America: New causal evidence. *Labour Economics, Elsevier*, 14(3):565–573, June 2007. URL <http://ideas.repec.org/a/eee/labeco/v14y2007i3p565-573.html>.
- Flavio Cunha and James Heckman. The technology of skill formation. *The American Economic Review*, 97(2):pp. 31–47, 2007. ISSN 00028282. URL <http://www.jstor.org/stable/30034418>.
- Flavio Cunha, James J. Heckman, and Susanne M. Schennach. Estimating the technology of cognitive and noncognitive skill formation. *Econometrica*, 78(3):883–931, 2010. URL <http://www.nber.org/papers/w15664>.



- Janet Currie. Healthy, wealthy, and wise: Socioeconomic status, poor health in childhood, and human capital development. *Journal of Economic Literature*, 47(1):87–122, March 2009. URL <http://ideas.repec.org/a/aea/jecclit/v47y2009i1p87-122.html>.
- Janet Currie and Rosemary Hyson. Is the impact of health shocks cushioned by socioeconomic status? the case of low birthweight. *The American Economic Review*, 89(2):pp. 245–250, 1999. ISSN 00028282. URL <http://www.jstor.org/stable/117114>.
- Janet Currie and Hannes Schwandt. Within-mother analysis of seasonal patterns in health at birth. *Proceedings of the National Academy of Sciences of the United States of America*, 110(30):1–6, 2013. doi: 10.1073/pnas.1307582110.
- Janet Currie and Tom Vogl. Early-life health and adult circumstance in developing countries. *Annual Review of Economics*, 5:1–36, 2013. doi: 10.1146/annurev-economics-081412-103704. URL <http://www.annualreviews.org/doi/abs/10.1146/annurev-economics-081412-103704>.
- Janet Currie and Reed Walker. Traffic congestion and infant health: evidence from e-zpass. *American Economic Journal: Applied Economics*, 3:65–90, 2011.
- Janet Currie and Aaron Yelowitz. Are public housing projects good for kids? *Journal of Public Economics*, 75:99–124, 2000.
- Janet Currie, Mark Stabile, Phongsack Manivong, and Leslie L Roos. Child health and young adult outcomes. *Journal of Human Resources*, 45:517–548, 2010.
- F.C. Curriero, J.A. Patz, J.B. Rose, and S. Lele. The association between extreme precipitation and waterborne disease outbreaks in the united states, 1948-1994. *American Journal of Public Health*, 91:1194–1199, 2001. URL <http://www.ncbi.nlm.nih.gov/pubmed/11499103>.

- Gordon B. Dahl and Enrico Moretti. The demand for sons. *Review of Economic Studies*, 75:1085–1120, 2008.
- J DaVanzo, L Hale, A Razzaque, and M Rahman. The effects of pregnancy spacing on infant and child mortality in matlab, bangladesh: how they vary by the type of pregnancy outcome that began the interval. *Population Studies*, 62(2):131–154, July 2008. doi: 10.1080/00324720802022089.
- Mercedes de Onis, Adelheid W Onyango, Elaine Borghi, Amani Siyam, Chizuru Nishida, and Jonathan Siekmann. Development of a who growth reference for school-aged children and adolescents. *Bulletin of the World Health Organization*, 85:660–667, 2007. URL [www.who.int/growthref/growthref\\_who\\_bull.pdf](http://www.who.int/growthref/growthref_who_bull.pdf).
- Arunachalam Dharmalingam. The social context of family size preferences and fertility behavior in a south indian village. *Genus*, 52(1/2):pp. 83–103, 1996. ISSN 00166987. URL <http://www.jstor.org/stable/29789222>.
- Esther Duflo. Schooling and labor market consequences of school construction in indonesia: Evidence from an unusual policy experiment. *The American Economic Review*, 91(4):pp. 795–813, 2001. ISSN 00028282. URL <http://www.jstor.org/stable/2677813>.
- Tim Dyson and Mick Moore. On kinship structure, female autonomy, and demographic behavior in india. *Population and Development Review*, 9(1):pp. 35–60, 1983. ISSN 00987921. URL <http://www.jstor.org/stable/1972894>.
- Avraham Ebenstein and Steven Leung. Son preference and access to social insurance: Evidence from china’s rural pension program. *Population and Development Review*, 36(1):47–70, 2010. ISSN 1728-4457. doi: 10.1111/j.1728-4457.2010.00317.x. URL <http://dx.doi.org/10.1111/j.1728-4457.2010.00317.x>.

- Michael D Eriksen and Stuart Rosenthal. Crowd out effects of place-based subsidized rental housing: new evidence from the lihtc pro. *Journal of Public Economics*, 94(11-12):953–966, 2010.
- J.G. Eriksson, T Forsen, J Tuomilehto, C Osmond, and DJ P Barker. Early growth and coronary hearth disease in later life: longitudinal study. *BMJ*, 322(7292):949–953, April 2001.
- Shelly Errington. Recasting sex, gender, and power: A theoretical and regional overview. In Jane Monnig Atkinson and Shelly Errington, editors, *Power and Difference: Gender in Island Southeast Asia*. Stanford University Press, 1990.
- Erica Field, Omar Robles, and Maximo Torero. Iodine deficiency and schooling attainment in tanzania. *American Economic Journal: Applied Economics*, 1(4):140–69, October 2009. URL <http://ideas.repec.org/a/aea/aejapp/v1y2009i4p140-69.html>.
- David Figlio, Sarah Hamersma, and Jeffrey Roth. Does prenatal wic participation improve birth outcomes? new evidence from florida. *Journal of Public Economics*, 93(1-2):235–245, February 2009. URL <http://www.sciencedirect.com/science/article/pii/S0047272708001266>.
- Steven M Fishman, Laura E Caulfield, Mercedes de Onis, Monika Blossner, Adnan A Hyder, Luke Mullany, and Rober E Black. Childhood and maternal underweight. In *Comparative Quantification of health risks: global and regional burden of disease attributable to selected major risk factors*. World Health Organization, 2004.
- Joseph Friedman and Daniel H Weinberg. *The economics of housing vouchers*. Academic Press, 1982.
- Joseph Friedman and Daniel H Weinberg. *The great housing experiment*. SAGE Publications, 1983.

- Ashok Gadgil. Drinking water in developing countries. *Annual Review of Energy and the Environment*, 23:253–286, 1998. doi: 10.1146/annurev.energy.23.1.253. URL <http://www.annualreviews.org/doi/pdf/10.1146/annurev.energy.23.1.253>.
- Paul Gertler, James Heckman, Rodrigo Pinto, Arianna Zanolini, Christel Vermeersch, Susan Walker, Susan M. Chang, and Sally Grantham-McGregor. Labor market returns to early childhood stimulation: A 20-year followup to an experimental intervention in jamaica. Technical report, NBER Working Paper, 2013.
- PD Gluckman and MA Hanson. Living with the past: evolution, development, and patterns of disease. *Science*, 305:1733–1736, 2004.
- Z Griliches. Sibling models and data in economics: beginnings of a survey. *Journal of Political Economy*, 87:S37–S64, 1979.
- D.J. Gubler, P. Reiter, K.L. Ebi, W. Yap, R. Nasci, and J.A. Patz. Climate variability and change in the united states: Potential impacts on vector- and rodent-borne diseases. *Environmental Health Perspectives*, 109:223–233, 2001. URL <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC1240669/>.
- Monica Das Gupta, Jiang Zhenghua, Li Bohua, Xie Zhenming, Woojin Chung, and Bae Hwa-Ok. Why is son preference so persistent in east and south asia? a cross-country study of china, india and the republic of korea. *The Journal of Development Studies*, 40(2):153–187, 2003.
- Hilary Hoynes, Marianne Page, and Ann Huff Stevens. Can targeted transfers improve birth outcomes? evidence from the introduction of the wic program. *Journal of Public Economics*, 95:813–827, 2011.
- Luoja Hu and Analia Schlosser. Prenatal sex selection and girls’ well-being: Evidence from india. IZA Discussion Papers 5562, Institute for the Study of Labor (IZA), March 2011. URL <http://ideas.repec.org/p/iza/izadps/dp5562.html>.

- Martin Huber. Identifying causal mechanisms in experiments (primarily) based on inverse probability weighting. Technical report, University of St. Gallen School of Economics Working Paper, 2012. URL <http://ideas.repec.org/p/usg/econwp/201213.html>.
- P.R. Hunter. *Waterborne Disease: Epidemiology and Ecology*. Wiley, 1997.
- Atsushi Inoue and Gary Solon. Two-Sample Instrumental Variables Estimators. *The Review of Economics and Statistics*, 92(3):557–561, August 2010. URL <http://ideas.repec.org/a/tpr/restat/v92y2010i3p557-561.html>.
- Brian A Jacob. Public housing, housing vouchers, and student achievement: evidence from public housing demolitions in chicago. *American Economic Review*, 94(1):233–258, 2004.
- Brian A Jacob and Jens Ludwig. The effects of housing assistance on labor supply: evidence from a voucher lottery. *American Economic Review*, 102(1):272–304, 2012.
- Seema Jayachandran. Selling labor low: Wage responses to productivity shocks in developing countries. *Journal of Political Economy*, 114(3):538–575, June 2006. URL <http://ideas.repec.org/a/ucp/jpolec/v114y2006i3p538-575.html>.
- Seema Jayachandran. The Roots of Gender Inequality in Developing Countries. NBER Working Papers 20380, National Bureau of Economic Research, Inc, August 2014. URL <http://ideas.repec.org/p/nbr/nberwo/20380.html>.
- Seema Jayachandran and Ilyana Kuziemko. Why do mothers breastfeed girls less than boys? evidence and implications for child health in india. *The Quarterly Journal of Economics*, 126(3):1485–1538, 2011. URL <http://ideas.repec.org/a/oup/qjecon/v126y2011i3p1485-1538.html>.

- Seema Jayachandran and Rohini Pande. Why are indian children shorter than african children? Technical report, Working Paper, 2013.
- Robert Jensen and Nolan Miller. Keepin' 'em down on the farm: migration and strategic investment in children's schooling. Technical report, Working Paper, 2011.
- Michael Kevane and David I. Levine. Changing status of daughters in indonesia. Technical report, Center for International and Development Economics Research, 2003. URL <http://escholarship.org/uc/item/0b52v28f>.
- Sunita Kishor. Gender differential in child mortality: A review of the evidence. In Monica Das Gupta, Lincoln C. Chen, and T.N. Krishnan, editors, *Women's Health in India: Risk and Vulnerability*. Bombay: Oxford University Press, 1995.
- Jeffrey R Kling, Jeffrey B Liebman, and Lawrence F Katz. Experimental analysis of neighborhood effects. *Econometrica*, 75(1):83–119, 01 2007. URL <http://ideas.repec.org/a/ecm/emetrp/v75y2007i1p83-119.html>.
- Gayatri B. Koolwal. Son preference and child labor in nepal: The household impact of sending girls to work. *World Development*, 35(5):881 – 903, 2007. ISSN 0305-750X. doi: <http://dx.doi.org/10.1016/j.worlddev.2007.01.001>. URL <http://www.sciencedirect.com/science/article/pii/S0305750X07000162>.
- Lori Kowaleski-Jones and Greg J. Duncan. Effects of participation in the wic program on birthweight: evidence from the national longitudinal survey of youth. special supplemental nutrition program for women, infants, and children. *Am J Public Health*, 92(5): 799–804, May 2002.
- M.S. Kramer, F. Aboud, E. Mironova, I. Vanilovich, R.W. Platt, L. Matush, S. Igumnov, E. Fombonne, N. Bogdanovich, T. Ducruet, J.P. Collet, B. Chalmers, E. Hodnett, S. Davidovsky, O. Skugarevsky, O. Trofimovich, L. Kozlova, and S. Shapiro. Breastfeeding and

- child cognitive development: New evidence from a large randomized trial. *Archives of General Psychiatry*, 65:578–584, 2008. URL <http://www.ncbi.nlm.nih.gov/pubmed/18458209>.
- David Levine and Michael Kevane. Are investments in daughters lower when daughters move away? evidence from indonesia. *World Development*, 31(6):1065–1084, 2003. ISSN 0305-750X. doi: [http://dx.doi.org/10.1016/S0305-750X\(03\)00050-0](http://dx.doi.org/10.1016/S0305-750X(03)00050-0). URL <http://www.sciencedirect.com/science/article/pii/S0305750X03000500>.
- David I. Levine and Dean Yang. A note on the impact of local rainfall on rice output in indonesian districts. Technical report, University of Michigan Working Paper, 2006.
- Lixing Li and Xiaoyu Wu. Gender of children, bargaining power, and intrahousehold resource allocation in china. *Journal of Human Resources*, 46(2):295–316, 2011. URL <http://ideas.repec.org/a/uwp/jhriss/v46y2011ii1p295-316.html>.
- Jeffrey M Lubell, Mark Shroder, and Barry Steffen. Work participation and length of stay in hud-assisted housing. *Cityscape: A Journal of Policy Development and Research*, 6(2):207–223, 2003.
- V.C. Lucia, Z. Luo, J.C. Gardiner, N. Paneth, and N. Breslau. Reports of birthweight by adolescents and their mothers: Comparing accuracy and identifying correlates. *Paediatric and Perinatal Epidemiology*, 20:520–527, 2006.
- R. Lukito. The enigma of national law in indonesia: the supreme court’s decisions on gender-neutral inheritance. *Journal of Legal Pluralism*, 52:147–167, 2006.
- Sharon Maccini and Dean Yang. Under the weather: Health, schooling, and economic consequences of early-life rainfall. *The American Economic Review*, 99(3):pp. 1006–1026, 2009. ISSN 00028282. URL <http://www.jstor.org/stable/25592491>.

- Karen Macours, Norbert Schady, and Renos Vakis. Cash transfers, behavioral changes, and cognitive development in early childhood: Evidence from a randomized experiment. *American Economic Journal: Applied Economics*, 4(2):247–73, April 2012. URL <http://ideas.repec.org/a/aea/aejapp/v4y2012i2p247-73.html>.
- John A. Maluccio, John Hoddinott, Jere R. Behrman, Reynaldo Martorell, Agnes R. Quisumbing, and Aryeh D. Stein. The impact of improving nutrition during early childhood on education among guatemalan adults. *Economic Journal*, 119(537):734–763, 04 2009. URL <http://ideas.repec.org/a/ecj/econjl/v119y2009i537p734-763.html>.
- Maurizio Mazzocco and Shiv Saini. Testing efficient risk sharing with heterogeneous risk preferences. *American Economic Review*, 102(1):428–68, February 2012. URL <http://ideas.repec.org/a/aea/aecrev/v102y2012i1p428-68.html>.
- Annamaria Milazzo. Why are adult women missing? Technical report, Policy Research Working Paper, 2014.
- Bryce Millett and Manisha Shah. The effects of in-utero shocks on cognitive test scores: Evidence from droughts in india. Technical report, Harvard University, 2012a.
- Bryce Millett and Manisha Shah. Could droughts improve human capital? evidence from india. Technical report, Harvard University Working Paper, 2012b.
- Gregory Mills, Daniel Gubits, Larry Orr, David Long, Judie Feins, Bulbul Kaul, Michelle Wodd, Amy Jones & Associates, Coudburst Consulting, and the QED Group. The effects of housing vouchers on welfare families. Technical report, U.S. Department of Housing and Urban Development, Office of Policy Development and Research, 2006.
- Jonathan Morduch. Income smoothing and consumption smoothing. *Journal of Economic Perspectives*, 9(3):103–114, Summer 1995. URL <http://ideas.repec.org/a/aea/jecper/v9y1995i3p103-14.html>.



- Oscar Newman. *Defensible space: crime prevention through urban design*. Macmillan, 1972.
- Sandra Newman and Joseph Harkness. Ahousing housing and the educational attainment of children. *Journal of Housing Economics*, 9:40–63, 2000.
- Philip Oldenburg. Sex ratio, son preference and violence in india: A research note. *Economic and Political Weekly*, 27(49/50):pp. 2657–2662, 1992. ISSN 00129976. URL <http://www.jstor.org/stable/4399215>.
- Edgar O Olsen, Catherine A Tyler, Jonathan W King, and Paul E Carrillo. The effects of different types of housing assistance on earnings and employment. *Cityscape: A Journal of Policy Development and Research*, 8(2):163–187, 2005.
- Philip Oreopoulos. The long-run consequences of living in a poor neighborhood. *Quarterly Journal of Economics*, 118(4):1533–1575, 2003.
- Owen Ozier. Long-term effects of early childhood deworming. Technical report, Development Economics Research Group Working Paper, 2011.
- Alberto Palloni and Marta Tienda. The effects of breastfeeding and pace of childbearing on mortality at early ages. *Demography*, 23(1):pp. 31–52, 1986. ISSN 00703370. URL <http://www.jstor.org/stable/2061406>.
- Rohini Pande and Nan Astone. Explaining son preference in rural india: the independent role of structural versus individual factors. *Population Research and Policy Review*, 26(1):1–29, February 2007. URL <http://ideas.repec.org/a/kap/poprpr/v26y2007i1p1-29.html>.
- Varsha R. Parchure. Bonds and options: Capital market solutions for crop insurance problems. Technical report, National Insurance Academy: Pune, 2002.

- DL Pelletier, EA Fronquillo Jr, DG Schroeder, and JP Habicht. The effects of malnutrition on child mortality in developing countries. *Bulletin of the World Health Organization*, 73:443–448, 1995.
- Mark M. Pitt and Wendy Sigle. Seasonality, weather shocks and the timing of births and child mortality in senegal. Technical report, Brown University Working Paper, 2012.
- Mark M. Pitt, Mark R. Rosenzweig, and Md. Nazmul Hassan. Productivity, health, and inequality in the intrahousehold distribution of food in low-income countries. *The American Economic Review*, 80(5):1139–1156, December 1990. URL <http://www.jstor.org/stable/2006766>.
- Mark M. Pitt, Mark R. Rosenzweig, and Mohammad Nazmul Hassan. Human capital investment and the gender division of labor in a brawn-based economy. *American Economic Review*, 102(7):3531–60, 2012. doi: 10.1257/aer.102.7.3531. URL <http://www.aeaweb.org/articles.php?doi=10.1257/aer.102.7.3531>.
- Nancy Qian. Missing women and the price of tea in china: The effect of sex-specific earnings on sex imbalance. *The Quarterly Journal of Economics*, 123:1251–1285, 2008. doi: 10.1162/qjec.2008.123.3.1251. URL <http://qje.oxfordjournals.org/content/123/3/1251.short>.
- Lupin Rahman and Vijayendra Rao. The determinants of gender equity in india: Examining dyson and moore’s thesis with new data. *Population and Development Review*, 30(2): 239–268, June 2004.
- Martin Ravallion and Shubham Chaudhuri. Risk and insurance in village india: Comment. *Econometrica*, 65(1):171–184, January 1997. URL <http://ideas.repec.org/a/ecm/emetrp/v65y1997i1p171-184.html>.
- Marie-Claire Robitaille. Determinants of stated son preference in india: Are men and women different? *The Journal of Development Studies*, 49(5):657–669, 2013.

- James E Rosenbaum. Changing the geography of opportunity by expanding residential choice: lessons from the gautreaux program. *Housing Policy Debate*, 6(1):231–269, 1995.
- Mark R Rosenzweig. Risk, implicit contracts and the family in rural areas of low-income countries. *Economic Journal*, 98(393):1148–70, December 1988. URL <http://ideas.repec.org/a/ecj/econjl/v98y1988i393p1148-70.html>.
- Mark R. Rosenzweig and T. Paul Schultz. Market opportunities, genetic endowments, and intrafamily resource distribution: Child survival in rural india. *The American Economic Review*, 72(4):803–815, September 1982. URL <http://www.jstor.org/stable/1810018>.
- Mark R Rosenzweig and Oded Stark. Consumption smoothing, migration, and marriage: Evidence from rural india. *Journal of Political Economy*, 97(4):905–26, August 1989. URL <http://ideas.repec.org/a/ucp/jpolec/v97y1989i4p905-26.html>.
- Mark R. Rosenzweig and Kenneth I. Wolpin. Heterogeneity, intrafamily distribution and child health. *The Journal of Human Resources*, 23:437–461, 1988. URL <http://www.jstor.org/stable/145808>.
- Mark R Rosenzweig and Kenneth I Wolpin. Credit market constraints, consumption smoothing, and the accumulation of durable production assets in low-income countries: Investment in bullocks in india. *Journal of Political Economy*, 101(2):223–44, April 1993. URL <http://ideas.repec.org/a/ucp/jpolec/v101y1993i2p223-44.html>.
- Heather Royer. Separated at girth: Estimating the long-run and intergenerational effects of birthweight using twins. *American Economic Journal: Applied Economics*, 1:49–85, 2009.

- Lisa Sanbonmatsu, Jens Ludwig, Lawrence F Katz, Lisa A Gennetian, Greg J Duncan, Ronald C Kessler, Emma Adam, Thomas W McDade, and Stacy Tessler Lindau. Moving to opportunity for fair housing demonstration program—final impacts evaluation. Technical report, U.S. Department of Housing and Urban Development, Office of Policy Development and Research, 2011.
- Mark Shroder. Does housing assistance perversely affect self-sufficiency? a review essay. *Journal of Housing Economics*, 11:381–417, 2002.
- David Stevens. Employment that is not covered by state unemployment insurance laws. Technical report, U.S. Census Bureau LEHD Technical Paper, 2007.
- James H. Stock and Motohiro Yogo. Testing for weak instruments in linear iv regression. NBER Technical Working Papers 0284, National Bureau of Economic Research, Inc, November 2002. URL <http://ideas.repec.org/p/nbr/nberte/0284.html>.
- Tahlim Sudaryanto, Sri Hery Susilowati, and Sony Sumaryanto. Increasing number of small farms in indonesia: Causes and consequences. Technical report, European Association of Agricultural Economists, 2009. URL <http://ageconsearch.umn.edu/handle/52808>.
- Scott Susin. Longitudinal outcomes of subsidized housing recipients in matched survey and administrative data. *Cityscape: A Journal of Policy Development and Research*, 8(2):189–218, 2005.
- Peter A Tatian and Christopher Snow. The effects of housing assistance on income, earnings, and employment. *Cityscape: A Journal of Policy Development and Research*, 8(2):135–161, 2005.
- Robert M Townsend. Risk and insurance in village india. *Econometrica*, 62(3):539–91, May 1994. URL <http://ideas.repec.org/a/ecm/emetrp/v62y1994i3p539-91.html>.

- Penelope Troude, Laurence Foix L'Helias, Anne-Marie Raison-Boulley, Christine Castel, Christine Pichon, Jean Bouyer, and Elise de La Rouchebrochard. Perinatal factors reported by mothers: Do they agree with medical records? *European Journal of Epidemiology*, 23:557–564, 2008.
- Cesar G Victoria, Linda Adair, Caroline Fall, Pedro C Hallal, Reynaldo Martorell, Linda Richter, and Harshpal Singh Sachdev. Maternal and child undernutrition: Consequences for adult health and human capital. *The Lancet*, 371:340–357, 2008. doi: 10.1016/S0140-6736(07)61692-4. URL <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2258311/>.
- Carol Vlassoff. The value of sons in an indian village: How widows see it. *Population Studies*, 44(1):pp. 5–20, 1990. ISSN 00324728. URL <http://www.jstor.org/stable/2174301>.
- S.P. Walker, S.M. Chang, M. Vera-Hernandez, and S. Grantham-McGregor. Early childhood stimulation benefits adult competence and reduces violent behavior. *Pediatrics*, 127(5):849–857, 2011.
- WHO. Who child growth standards. Technical report, The World Health Organization, 2006. URL [www.who.int/childgrowth/standards/Technical\\_report.pdf](http://www.who.int/childgrowth/standards/Technical_report.pdf).
- WHO. Global database on child growth and malnutrition. web site: <http://www.who.int/nutgrowthdb/about/introduction/en/index2.html> [last accessed: 25 September 2014], 2014.
- Guoyao Wu, Fuller W Bazer, Timothy A Cudd, Cynthia J Meininger, and Thomas E Spencer. Maternal nutrition and fetal development. *Journal of Nutrition*, 134:2169–2172, 2004.
- Kazuo Yamaguchi. A formal theory for male-preferring stopping rules of childbearing:

Sex differences in birth order and in the number of siblings. *Demography*, 26(3):pp. 451–465, 1989. ISSN 00703370. URL <http://www.jstor.org/stable/2061604>.

Kathryn M. Yount. Women's family power and gender preference in minya, egypt. *Journal of Marriage and Family*, 67(2):pp. 410–428, 2005. ISSN 00222445. URL <http://www.jstor.org/stable/3600278>.