Journal of Development Economics 126 (2017) 52–65

Contents lists available at ScienceDirect

Journal of Development Economics

journal homepage: www.elsevier.com/locate/jdeveco

Extended families and child well-being☆

Daniel LaFave[a,*,1], Duncan Thomas[b,2]

[a] Colby College, United States
[b] Duke University, United States

ARTICLE INFO

Keywords:
Extended families
Pareto efficiency
Collective rationality
Child development

ABSTRACT

Whereas the extended family plays a central role in many models of economic behavior, particularly in developing countries, there is a paucity of empirical evidence on the extent and nature of resource sharing among non coresident family members. This is in sharp contrast with abundant evidence that the distribution of resources within households predicts household spending and savings patterns. To fill this gap, the collective model of household decision-making is extended to the family. The model is particularly appealing in this context because it places few restrictions on preferences of individual family members who may or may not be coresident and does not specify a specific bargaining mechanism that underlies negotiations. The model yields empirical tests of whether the behavior of family members is (Pareto) efficient.

Evidence is presented on the relationship between three distinct measures of health- and education-related human capital of children and the distribution of wealth among extended family members using rich longitudinal survey data from the Indonesia Family Life Survey (IFLS). The data are ideally suited for this research because the survey follows family members when they leave the family home and detailed information about individual-specific wealth is collected. We find that child human capital outcomes are affected by wealth of non coresident family members indicating that extended families do share resources. While the special case of the model in which all members are completely altruistic is rejected, the restrictions of the efficient model are not rejected, indicating that non co-resident family members are able to co-ordinate allocation decisions in such a way as to make no family member better off without another member being worse off.

1. Introduction

An active line of inquiry highlights the role of membership in social networks, villages and families to enable sharing of resources, risks and information in order to smooth consumption and invest in human capital. These linkages are especially important in low income settings where markets for insurance and credit are incomplete, where publicly-provided safety nets are limited and where social security is not widespread. (See, for example, Angelucci and de Giorgi, 2009; Angelucci et al., 2010; Cox and Jimenez, 1990; Fafchamps and Lund, 2003; Jensen, 2003; Munshi, 2003; Rosenzweig and Stark, 1989; Smith and Thomas, 1998; Townsend, 1994). In much of this work, the extended family is thought to play a central role in providing resources to members in times of need. There is, however, a paucity of evidence on the extent to which families share resources and the impact of that sharing on the well-being of family members. Given the importance of the family in many models in economic behavior, this is an important gap in the literature.

In sharp contrast, a large number of studies has established that variaton in the distribution of resources among coresident household members is predictive of variation in spending and savings patterns, with females typically allocating more resources to investments in the future – including their children – relative to resources allocated by males (Thomas, 1990; Lundberg, Pollak, and Wales, 1997; Rangel, 2006; Ashraf, 2009; Rubalcava, Teruel and Thomas, 2009; Bobonis, 2009). It is, however, not straightforward to draw conclusions from evidence on intra-household resource allocation about how variation in the distribution of resources among non coresident family members affects resource allocation decisions in the family. For example, in South Africa, children in households that have a member who is eligible for the old age pension tend to be taller (Duflo, 2003). However, prime-age adults, who are likely to be the parents of those children, tend to be

☆ We are grateful to Joseph Hotz, Rob Jensen, Marjorie McElory, Robert Pollak, Marcos Rangel, David Ribar, Seth Sanders, Alessandro Tarozzi, and numerous seminar participants for many helpful suggestions and comments.

* Corresponding author.
E-mail addresses: daniel.lafave@colby.edu (D. LaFave), dthomas@econ.duke.edu (D. Thomas).
[1] Department of Economics, Colby College, Waterville, ME 04901, United States.
[2] Department of Economics, Duke University, Durham, NC 27708, United States.

http://dx.doi.org/10.1016/j.jdeveco.2016.11.006
Received 8 July 2015; Received in revised form 11 July 2016; Accepted 28 November 2016
Available online 08 December 2016
0304-3878/

taller and better educated which cannot be attributed to the impact of the pension income but likely reflects shifts in living arrangements as higher human capital children and grandchildren move in when the grand-parent becomes old enough to receive the pension (Hamoudi and Thomas, 2014). This example highlights the importance of taking into account endogenous living arrangements in order to draw inferences about the impact of resources in the hands of one family member on well-being of all members.

The goal of this paper is to reach beyond the allocation of resources among coresident household members to investigate the relationship between the well-being of individuals within an extended family and the distribution of resources among coresident and non coresident family members. Focusing on human capital outcomes of children, we define an extended family as individuals who are biologically linked to the child including the child's parents, grandparents and siblings as well as the siblings of the child's parents.

Using data from the Indonesia Family Life Survey (IFLS) we examine three markers of human capital of young children that are thought to be affected by resources (Heckman, 2006). They are, first, height, conditional on age and gender, a measure of health and nutrition; second, performance on a non-verbal cognitive assessment and, third, age the child started school.

The theoretical foundation for our empirical tests is an extension of the collective model of household decision-making (Chiappori, 1988, 1992) to the context of decisions made by non co-resident family members. The extended collective model is well-suited to the application: it makes no assumptions about co-residence choices, does not impose a specific structure for how family members bargain with one another, and yields empirical tests about the nature of resource sharing within the family.

The key assumption in the model is that allocations within the extended family are Pareto efficient in the sense that no family member can be made better off without another family member being made worse off. In the context of a household, failure of efficiency may be difficult to rationalize since living arrangements are properly treated as endogenous. However, in the context of family decision-making, it is not obvious that non co-resident family members will achieve efficient allocations because of co-ordination costs, asymmetric information or because of preferences. It is, therefore, a substantively interesting test of family behavior that places plausible, testable restrictions on the behavior of individual family members who likely have heterogeneous preferences.

Indonesia provides an ideal context for research on the extended family. First, the anthropological literature discusses inter-linkages between kin to provide resources not only when a family member experiences a negative shock, but also to support investments in the next generation and to provide aid at older ages. This literature has emphasized the importance of trust and reciprocity within kinship networks drawing on qualitative evidence from villages in Indonesia (Geertz, 1961, 1963; Jay, 1969; Schroder-Butterfill, 2005; Kreager and Schroder-Butterfill, 2007).

Second, studies have extensively described the importance of transfers among non co-resident family members in Indonesia. Frankenberg et al. (2002) document the fact that exchanges between non co-resident family members are widespread and conclude that there is evidence in support of at least three motivations for these exchanges. First, some transfers appear to reflect an exchange of money for time as parents buy filial obligations (Park, 2003). Second, it has been suggested that parents pay for their children's education partly as a loan that is later repaid in the form of old age support (Raut and Tran, 2005; Cameron and Cobb-Clark, 2008). Third, transfers reflect insurance among family members. Park (2003) also documents that older siblings assist younger siblings who have fewer resources. Okten and Osili (2004) describe information sharing among non-coresident siblings. While this work documents extensive financial links between non-coresident family members in Indonesia, it is difficult to establish a direct link between transfers and specific spending outcomes with the survey data used in these studies.

Third, a small number of studies have investigated risk-sharing among non co-resident family members in Indonesia. Whereas Gertler and Gruber (2002) find imperfect consumption insurance against health shocks, a re-analysis by Genoni (2012) using more recent, richer data establishes that consumption insurance is almost complete and driven by resources from extended family members. The key role of risk-sharing among non-coresident family members in the face of the Asian financial crisis and the very broad set of mechanisms used to smooth consumption are described in detail in Frankenberg et al. (2003) and Thomas and Frankenberg (2007). Witoelar (2013) documents characteristics of the kin network that affect household consumption patterns in research that is closely related to the research reported here.

This paper makes four contributions to the literature. First, studies of household decision-making are, by design, limited to behavior of the self-selected group of individuals who have chosen to co-reside in a household unit. By placing the family at the center of decisions about sharing of resources, our paper directly addresses this important gap in the literature. The lack of evidence on how families make decisions is primarily a reflection of a paucity of data. Since household surveys are designed to collect information on members of sampled households, it is only with long-term panel studies which successfully track movers that these questions can be adequately addressed. Studies like the Panel Study of Income Dynamics and IFLS are especially well-suited data sources for these types of analyses.

Second, seminal work by Altonji et al. (1992) and Hayashi et al. (1996) test whether families are completely altruistic and behave as if all members share the same preferences (or one family member makes all decisions). Using expenditure data from the U.S., they reject this unitary model of the family. We test a more general model of resource-sharing that tests whether allocations are Pareto efficient. Along the way, we also test, and reject, the special case of unitary families. The results of our test for efficient allocations are of substantive interest because they determine whether non coresident family members are able to co-ordinate actions so that resources are shared in such a way that no family member could be made better off without another member being worse off.

Third, while examinations of how households allocate their budgets are a powerful tool for making welfare comparisons, it is not clear how to interpret evidence of resource allocation of non co-resident family members that is based on expenditure patterns as in Altonji et al. (1992) and Witoelar (2013). Those studies compare the impact of a household's own per capita expenditure and the per capita expenditure of non co-resident family members on a household's spending on specific goods. Shared housing costs and economies of scale of living arrangements, including shared food purchase and preparation costs, are key mechanisms through which family members share resources. Thus, the distribution of expenditure across households in the family likely reflects the outcome of resource sharing decision and so is properly treated as endogenous in these models (Dalton et al., 2016). These issues also complicate interpretation of specific expenditures or budget shares. We address this limitation by exploiting uniquely-rich data collected in IFLS on the distribution of wealth among individual family members and side-step the problem of the endogeneity of household expenditure.

Fourth, rather than focus on budget shares, we examine child human capital outcomes which are likely to be affected by resources and have an immediate welfare interpretation. Whereas there is indirect evidence that child human capital outcomes are affected by resources outside the household, direct evidence on this question is scarce. For example, creative research has shown that eligibility for Progresa, the Mexican anti-poverty program, is associated with increased rates of school enrollment among children who share the same surname but live in a different household in the village (Angelucci et al.,

2010). Other studies suggest that resources of extended family members living outside the village, including those living in the United States, influence health and education outcomes of children in Mexico (Antman, 2011; Nobles, 2013; Farfan et al., 2016).

We show that child human capital outcomes are affected not only by the resources under control of their parents, but also resources of extended family members. We reject the unitary model of family behavior, or complete altruism among family members, but cannot rule out that family members co-ordinate actions in such a way as to allocate resources that affect child human capital outcomes in a manner that is consistent with Pareto efficiency. While the study is not designed to identify specific mechanisms or precise motivations that underlie this form of sharing, the results do have important implications for understanding several key issues in development. These include, for example, the intergenerational transmission of human capital, the evolution of inequality within a society and the likely efficacy of public programs targeted at specific individuals or households.

2. Theoretical and empirical framework

It is, in principle, straightforward to extend the model of household behavior of Chiappori (1988, 1992) and Browning et al. (1994) with home production (Chiappori, 1997) to the family context. The welfare of an extended family is maximized subject to a family level budget constraint that includes the resources of all family members. Welfare of the family is assumed to depend on the utility of each of its members,and, in turn, each member's sub-utility is assumed to depend on their own consumption of goods, services and leisure as well as that of all other extended family members. In addition, and key for this research, each member's sub-utility depends on the output of home production by the family which includes different domains of health, education and cognitive development. We focus on the human capital of children which depends on time and goods inputs as well as endowments, with output of each family member being constrained by the technology of the human capital production functions.

Examining child human capital outcomes is of substantial interest in models of family behavior for two reasons. First, non co-resident family members have been shown to invest in many dimensions of the health and well-being of the next generation. Second, there are multiple distinct dimensions of child human capital that are readily measured in a survey setting. While the model does not specify the functional form of the sub-utility functions, it is necessary for each sub-utility function to be quasi-concave, non-decreasing, and strictly increasing in at least one argument.

The sub-utility of each family member depends on individual, household, and family specific characteristics some of which are observed, such as age, household demographic structure, anthropometrics, and socioeconomic status, and some of which are not observed, such as attitudes toward human capital investments, altruism towards family members and human capital endowments. The extended family welfare function is a weighted sum of individual sub-utilities with the weight associated with each family member depending on the characteristics of both that individual and all other family members. Some of these characteristics are observed and include, for example, marriage market opportunities, prices (including the price of labor), age, education and the economic resources of each family member. Other characteristics, which are not observed, include, for example, time preferences and altruism. Intuitively, these weighting functions reflect the influence of members in the family decision-making process and can be interpreted as indicative of each member's bargaining power. A key advantage of this model over many of the co-operative and non co-operative bargaining models that have been discussed in the literature is that there is no need to specify an equilibrium concept for the bargaining process (see Chiappori (1988), as well as Lundberg and Pollak (1996), and Chiappori and Mazzocco (2016), for reviews).

2.1. Unitary model

Family decisions will be unitary if all family members are completely altruistic, if all family members have the same preferences or if the preferences of only one family member (a dictator) determines resource allocations (Samuelson, 1956; Becker, 1981). Under these conditions, child human capital outcomes depend only on total extended family resources and not on its distribution within the family. That is, conditional on total family resources, child human capital outcomes will not depend on the distribution of those resources between the child's parents, grandparents, aunts, uncles or other relatives. While this model provides a useful starting place for investigating family decision-making, it has been rejected in a large number of studies of household decision-making. It is, therefore, unlikely to be a satisfactory model of decision-making that includes non co-resident family members and evidence to this effect is provided below. We turn, therefore, to the substantially more general collective model.

2.2. Collective model

The collective model allows both total family resources and the distribution of resources within the family to affect resource allocations, but restricts allocations to be Pareto efficient. This is not an innocuous restriction, particularly in the context of decision-making by non co-resident family members. Results from empirically testing this restriction provides important insights into understanding family behavior.

Chiappori (1988, 1992) and Browning and Chiappori (1998) establish that the collective model is formally equivalent to a model of family income sharing in which decisions are made in two stages. First, all resources are pooled and then shared among family decision-makers. Second, each decision-maker allocates his or her share of the pooled resources to maximize own sub-utility without reference to allocations by other family members.

Each decision-maker's share of pooled resources is determined by a sharing rule that depends on the utility weights that aggregate sub-utility functions or, intuitively, on factors that affect the bargaining power of that family member in allocation decisions such as command over current and future resources, marriage market opportunities and altruism.

In the second stage, each family decision-maker maximizes his or her own utility subject to the human capital production functions and an individual-specific budget constraint that limits that member's spending to be no more than the member's share of resources from the first stage allocation mechanism, R_m. Member m's share of resources, R_m, depends on assets of each family member, $y_1,\ldots, y_M$, prices, p, along with observed, π, and unobserved, ξ, characteristics that affect the utility weights so the human capital outcome k of child i, θ_{ki}, is:

$$\theta_{ki} = \theta_{ki}(R_m(y_1, \ldots, y_M, p, \pi, \xi), p, \mu, \varepsilon) \quad (1)$$

which also depends on prices and observed, μ, and unobserved, ε, factors that affect the shape of the human capital production function and the family welfare function. Some, but not necessarily all, of these factors may also affect the sharing rule.

Whereas in the unitary model, allocations depend only on total family resources, in the collective model, the resources of each decision-maker affect allocations and so the distribution of resources within the family matters. The weakly separable sharing rule function (1) plays a key role in deriving empirical tests of the collective model (Bourguignon et al., 1993; Browning and Chiappori, 1998).

The marginal impact of assets in the hands of family member m, y_m, on demand for θ_{ki} is:

$$\frac{\partial \theta_{ki}}{\partial y_m} = \frac{\partial \theta_{ki}}{\partial R}\frac{\partial R}{\partial y_m} \tag{2}$$

where the first term on the right hand side reflects the impact of m's share of family resources from the first stage on human capital which, because of weak separability in (1) is the same for all family members no matter what drives variation in each members share from the first stage family allocation decision. This is the key empirical implication of family allocations being assumed to be Pareto efficient and lies at the heart of the tests of efficiency described below. The impact of member m's assets, y_m on the sharing rule is captured in the second term. The marginal impact of member n's assets, y_n, on the same human capital outcome, θ_{ki}, is:

$$\frac{\partial \theta_{ki}}{\partial y_n} = \frac{\partial \theta_{ki}}{\partial R}\frac{\partial R}{\partial y_n} \tag{3}$$

where the first term is the family income effect and thus the same as the first term in (2). Thus, if resource allocations are efficient, the distribution of assets within the family affects human capital outcomes only through their impact on the first stage sharing rule and so the ratio of the marginal effects of assets y_{m} and y_{n} is independent of θ_{ki} for all human capital outcomes:

$$\frac{\frac{\partial \theta_{ki}}{\partial y_m}}{\frac{\partial \theta_{ki}}{\partial y_n}} = \frac{\frac{\partial \theta_{ki}}{\partial R}\frac{\partial R}{\partial y_m}}{\frac{\partial \theta_{ki}}{\partial R}\frac{\partial R}{\partial y_n}} = \frac{\frac{\partial R}{\partial y_m}}{\frac{\partial R}{\partial y_n}} \quad \forall\ k,\ m \text{ and } n,\ m \neq n \tag{4}$$

for all pairs of family decision-makers, m and n. Put another way, the ratio of the marginal effects of each pair of family member's assets is constant across all human capital outcomes.

From an empirical perspective, this is a powerful result for at least two reasons. First, there is no a priori reason to expect the ratio of marginal propensities to be the same across all goods for a pair of family members. Second, it is straightforward to test (4) with survey data. The next sub-section describes the empirical implementation of the implications of the theoretical models described above.

2.3. Empirical implementation

Collecting terms in (1) and taking a linear approximation, the demand for θ^k_{ihf}, the human capital outcome k of child i in household h and family f, depends on y_m, assets of each family member, m, some of whom are members of households h and some are not, in addition to a set of demographic controls for the index child, X_{if}, and all other family members, X_{mf}:

$$\theta^k_{ihf} = \alpha^k + y_{mf}\beta^k_m + X_{if}\gamma^k_i + X_{mf}\gamma^k_m + \varepsilon^k_{if} \tag{5}$$

Exploiting the fact that IFLS collects information on resources at the individual level, we estimate (5) identifying individual family members, m, who are likely to be the most salient to the index child: the father, mother, grandfather, grandmother and other family members (which includes the child's siblings, aunts and uncles). From (4), allocations are collectively rational if the ratio of the marginal effect of assets of one family member, m, to the marginal effect of assets of another family member, n, is the same for all pairs of child outcomes, j and k:

$$\frac{\beta^k_m}{\beta^k_n} = \frac{\beta^j_m}{\beta^j_n} \quad \forall \quad j \neq k \tag{6}$$

In some cases, the value of an asset is well-known and assignment to an individual is straightforward; gold is a good example. Its value is a function of weight, the price per gram is set in world terms and gold is typically owned by an individual in a family. In other cases, respondents may have difficulty valuing the asset and apportioning ownership among family members. Valuations are most difficult for those assets that are seldom traded in the market, such as land. Apportioning the value of these assets is likely to be most difficult within couples and possibly within households.

In the empirical analyses, these concerns are addressed in two ways. First, we separate liquid assets (gold, cash and stocks, for example) from illiquid assets (land and housing). Second, we begin by putting aside dynamics within households and compare the impact of assets of those family members who live in the same household, h, as an index child with assets of those who do not co-reside, $\bar{h}$, with the child:

$$\theta^k_{ihf} = \alpha^k + y_{hf}\beta^k_h + y_{\bar{h}f}\beta^k_{\bar{h}} + X_{if}\gamma^k_i + X_{mf}\gamma^k_m + \varepsilon^k_{if} \tag{7}$$

In this special case of (5), allocations between co-resident and non co-resident family members will be collectively rational if the ratio of asset effects are the same for each pair of child outcomes:

$$\frac{\beta^k_h}{\beta^k_{\bar{h}}} = \frac{\beta^j_h}{\beta^j_{\bar{h}}} \quad \forall \quad j \neq k \tag{8}$$

Moreover, families behave as if they are unitary if child outcomes are invariant to whether resources are in the hands of household members or in the hands of family members outside the household, implying that the asset effects would then be equal, $\beta^k_h = \beta^k_{\bar{h}}$. A test that is robust to measurement error in total family assets replaces $y_{\bar{h}f}$ in (7) with a family-specific fixed effect, μ^k_f

$$\theta^k_{ihf} = \alpha^k + y_{hf}\beta^k_h + X_{if}\gamma^k_i + X_{mf}\gamma^k_m + \mu^k_f + \varepsilon^k_{if} \tag{9}$$

which effectively controls total family resources. In this case, household resources should have no additional impact on child outcomes and family behavior is unitary if $\beta^k_h = 0$.

Since the test for collective rationality involves cross equation restrictions, the child outcome models are estimated jointly to construct nonlinear Wald test statistics that are specified in terms of the products of coefficients in (6) and (8) by cross-multiplying each side of the equation. This formulation assures neither term in the equality is divided by zero and thus performs better than testing the ratios themselves. All variance-covariances are estimated allowing for intra-family correlation at the family level.[3] This is potentially important as we are comparing human capital outcomes of children in the same family who not only share genes but also have shared environments.

3. Data

Data are drawn from IFLS, a large-scale, ongoing longitudinal survey that collects detailed information on individuals, households and families. The 1993 baseline enumerated about 22,000 members of 7,224 households in 306 rural or urban communities (*desa*) spread across 13 of the Indonesian provinces. Individuals were re-interviewed in IFLS2 in 1997 and in IFLS3 in 2000. We use data collected in IFLS4, conducted in late 2007 and early 2008, which interviewed over 43,500 individuals living in 13,500 households in 3,650 *desa*.[4] The fact that the number of communities in which respondents were interviewed increases by over 12-fold between baseline and IFLS4 follow-up 15 years later reflects a key strength of IFLS for this research: baseline respondents and their children born after baseline who move are tracked and interviewed in their new location. Specifically, as children age, strike out on their own and form new households, they, their spouses and their children are interviewed. In IFLS4, 90.6% of all surviving IFLS respondents were re-interviewed (Thomas et al., 2012). This far exceeds the recontact rates in any other large-scale population-based survey including, for example, the Panel Study of Income Dynamics in which attrition was around 50% after 15 years.

[3] Inferences do not change if they are based on bootstrapped estimates that are clustered at the family level. Inference are no different if clustering is at the desa, kabupaten or province levels.

[4] See Frankenberg and Karoly (1995), Frankenberg and Thomas (2000) and Strauss et al. (2004, 2009) for descriptions of the waves of IFLS.

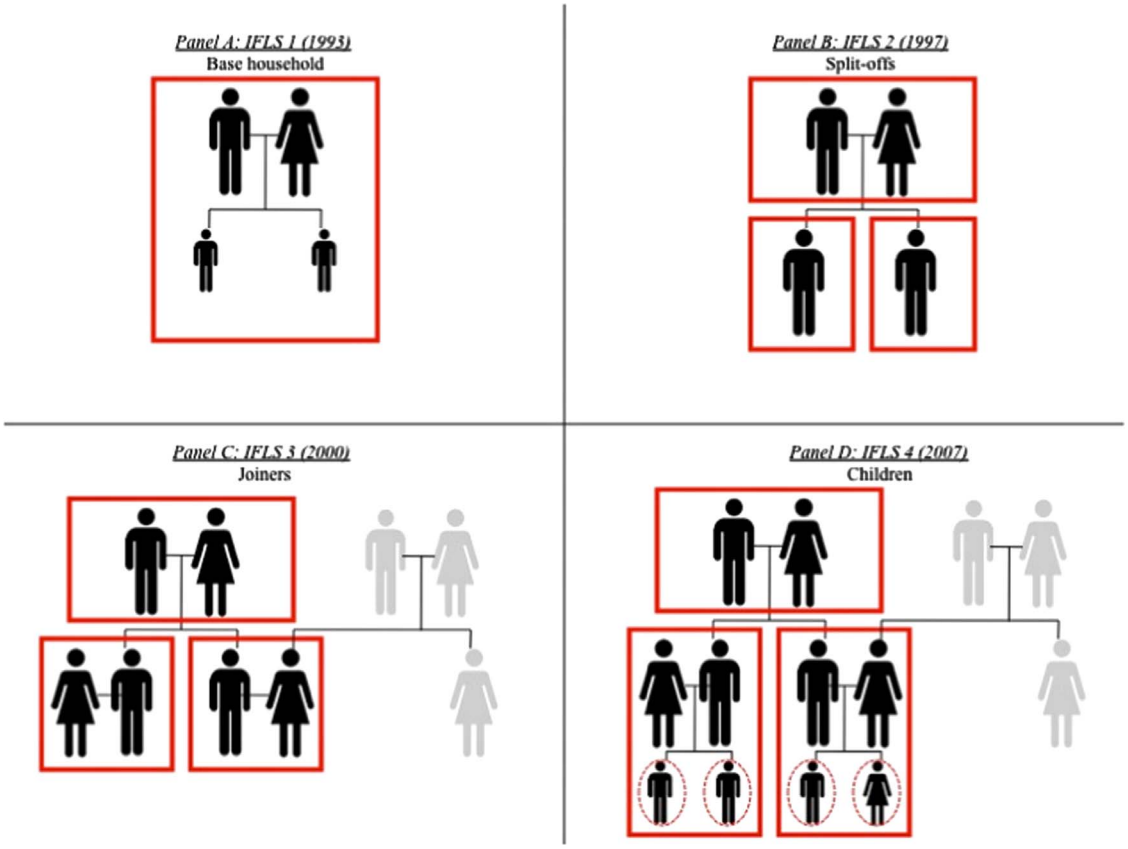

Fig. 1. Constructing Extended Families in the Indonesia Family Life Survey. *Notes*: This figure depicts how extended families are identified in the IFLS. Each box represents an IFLS household. Starting from a baseline household in the first wave of IFLS in 1993, when children split-off and form their own households they are followed and become part of the IFLS sample, as shown in Panel B. When spouses join the newly formed households, as in Panel C, we obtain information on the spouses and their new non-IFLS relatives, but only the spouse is part of the IFLS sample. Our analytical sample consists of the young children shown in Panel D and their families in IFLS4.

Following Altonji et al. (1992), we construct extended families that include non co-residents by linking individuals in IFLS4 who have split off from their baseline households with all IFLS respondents that have been spawned by the baseline household.

An example is illustrated in Fig. 1. Consider a baseline household with two parents and their two children. Four years later in IFLS2, the two children have grown and split-off to form their own households. These households are located and interviewed and we construct a family consisting of three households. In IFLS3, the baseline children have married, and their spouses become part of the IFLS sample. The spouses are administered the same individual interview as the original panel members. The light gray individuals indicate parents and siblings of joiners about whom information is collected from the joiner regarding vital status, age, gender education and transfers. As shown in Fig. 1, by IFLS4, each of the original children now has two children. Our empirical models examine the outcomes of these four children and relate those outcomes to the resources of their parents, their grandparents and other non co-resident family members.

In a household-based survey, the full genealogical family tree is not observed; the implications for our analyses are discussed in detail below. The key point for this research is that, by design, IFLS enables the construction of families of related kin who are not co-resident. Appendix Table A1 reports the number of children, parents and grand-parents as well as households and families included in our sample. The family structure of our sample is rich and varied. There are, on average, 8.7 members in a family, 4.5 members in a household and 2.3 households in a family. There are 1.7 children age 16 and under in the average household and half the households include a prime age adult who is not the parent of the child, with one-third of households including a grand-parent. Three-quarters of families are three generational and families are spatially-diversified with half having members who live in different districts and one in five having members who live in different provinces. In all, there are approximately 15,000 age-eligible children (age 0–16 yrs) in the samples used in this research.[5]

Child Outcomes The analyses examine three distinct indicators of child health and human capital that are related to well-being in later life: height-for-age, cognition, and participation in early schooling (kindergarten). Height has been established as an informative indicator of nutritional status during very early childhood that reflects the combination of genetic endowment and the influence of the disease and nutrition environment during the first two to three years of life (Waterlow et al., 1977; Martorell and Habicht, 1986). Studies have shown that resources such as parental time and family resources are predictive of child linear growth primarily through their influence on food intake and disease insults (Silventoinen et al., 2000; Alderman et al., 2006; Hoddinott et al., 2013). Moreover, child height is a powerful predictor of attained height as an adult which is associated with reduced mortality and morbidity as well as greater economic prosperity (Fogel, 2004; Strauss and Thomas 1998). Height or length of each child is measured by a trained enumerator and standardized for

[5] Taking all panel respondents who would be aged between 20 and 50 at the time of the survey (and thus potentially a parent of an age-eligible child), overall attrition is 22%. Importantly, attrition is very similar for males and females in this age group (the difference is 2.6%) and declines with the size of the household at baseline (by 1.8% for each additional household member) so that attrition is lower for larger extended families.

age and gender using the 2000 Center for Diseases Control growth tables which are normed to a representative, well-nourished child in the United States.[6] Since child linear growth is most sensitive to variation in inputs during the first few years of life (Martorell and Habicht, 1986), we focus on standardized height-for-age of children six years old and younger expressed as a z-score. Height is measued for 6,567 children which is 99.2% of all age-eligible children in the survey. A z-score of 0 indicates the child is the same height as a well-nourished U.S. child of the same age and gender. As shown in the second column of Table 1, the average child in our sample is slightly over 1 standard deviation shorter than the U.S. norm.

The second child outcome is a measure of cognitive function. It comprises performance on a non-verbal cognitive battery, the Raven's Colored Progressive Matrices (CPM) pattern recognition assessment and answers to five simple arithmetic questions that involve addition, subtraction, multiplication and division. The CPM assessment has been interpreted as a measure of fluid intelligence (Heckman and Kautz, 2012), the rate at which an individual learns, and a measure of Spearman's general intelligence factor *g* (e.g. Raven, 2000; Kaplan and Saccuzzo, 1997). Each child 7 and older completed a booklet with 12 question from the CPM battery. We combine the CPM and arithmetic questions and use the percentage of questions answered correctly for children age 7 through 16. The Raven's assessment was completed by 96.5% of all children in this age range in the survey and there are 7,727 children in our analytical sample. We flexibly control for age in all models as older children are more likely to answer correctly than younger children. The average score is 69.6% correct with the average Ravens score being 74.6% and the average arithmetic score being 61.4%.

Whereas height and cognition are outputs of investments by family members, our third indicator is an input into the education process, whether or not the child participated in kindergarten. Only about half the sample children attended kindergarten although school readiness is thought to be an important contributor to learning in primary school. More generally, early education programs have been shown in rigorous studies to have a positive impact on socio-emotional development, school readiness and, in some studies, on school attainment and later life outcomes (Knudsen et al., 2006; Gertler et al., 2014). In our data, attending kindergarten is positively and significantly associated with subsequent school outcomes, even after controlling parental education and resources.

These markers and the age ranges for each marker are selected because they capture different domains of human capital that are infuenced by different types of resources at different points in the life course. Whether a child attains their potential in physical stature is largely determined by nutrition and disease exposures during the first few years of life. Cognitive achievement reflects the roles of intellectual stimulation, reading and language development throughout childhood. Participation in kindergarten is more closely tied to resource availability at age 5 or 6.

Resources of family members The distribution of resources within the family plays a central role in the theory and empirical models. Child outcomes are invariant to this distribution in the unitary model, but are influenced by the distribution among family members through intra-family bargaining in the collective model. It is, however, a substantial challenge to translate the theory to an empirical specification of resources that reflects an individual's capacity to assert his or her own preferences over allocation decisions. Intuitively, one might think of resources as being indicative of an individual's bargaining power within a family say, for example, because the individual would retain control over those resources if all links with the family were severed.

[6] The choice of standards is immaterial; estimates based on World Health Organization and Center for Diseases Control standards are not different and no inferences are different.

Table 1
Descriptive statistics.

	Sample			
	All Children (Birth – 16 yrs)	Height-for-Age (Birth – 6 yrs)	Cognitive Score (7–16 yrs)	Attended Kindergarten (6–14 yrs)
Outcomes				
Height-for-Age (z-score)		−1.14		
		(0.02)		
Cognitive Score (%)			69.6	
			(0.21)	
Attended Kindergarten (%)				50.8
				(0.58)
*Illiquid Assets of [...]**				
Household	6660.5	6376.0	6946.2	6801.4
	(110.9)	(165.8)	(151.3)	(151.4)
Family	9431.8	11312.4	7845.2	8774.0
	(184.8)	(297.5)	(231.8)	(256.1)
*Liquid Assets of [...]**				
Household	1928.8	1949.1	1902.7	1854.7
	(38.9)	(60.3)	(51.5)	(49.7)
Family	2579.5	3027.9	2148.9	2355.9
	(47.3)	(74.6)	(59.5)	(64.7)
Mother	762.8	680.6	831.3	822.4
	(17.7)	(22.1)	(27.1)	(26.8)
Father	872.2	821.5	904.2	907.9
	(22.9)	(37.1)	(28.2)	(29.9)
Grandmother	616.7	697.4	480.5	510.8
	(19.9)	(29.4)	(23.7)	(22.7)
Grandfather	799.8	885.6	657.9	658.3
	(35.3)	(51.5)	(42.2)	(37.7)
Additional Controls				
Age	7.37	2.88	11.26	9.79
	(0.04)	(0.02)	(0.03)	(0.03)
Female	0.48	0.48	0.49	0.48
	(0.00)	(0.01)	(0.01)	(0.01)
Co-reside with Mother	0.90	0.96	0.85	0.87
	(0.00)	(0.00)	(0.00)	(0.00)
Co-reside with Father	0.81	0.87	0.76	0.78
	(0.00)	(0.00)	(0.00)	(0.00)
Household Size	4.98	4.86	5.08	5.06
	(0.02)	(0.02)	(0.02)	(0.02)
Family Size	11.16	11.94	10.45	10.86
	(0.05)	(0.08)	(0.07)	(0.07)
Age of [...]				
Mother	34.1	29.8	37.9	36.4
	(0.07)	(0.08)	(0.08)	(0.08)
Father	38.7	34.4	42.5	41.0
	(0.08)	(0.09)	(0.10)	(0.10)
Grandmother	58.5	55.9	62.1	61.2
	(0.11)	(0.14)	(0.17)	(0.16)
Grandfather	62.0	59.9	65.5	64.7
	(0.13)	(0.16)	(0.21)	(0.19)
Years of Education of [...]				
Mother	8.03	8.96	7.26	7.56
	(0.03)	(0.05)	(0.05)	(0.05)
Father	8.52	9.15	7.96	8.22
	(0.04)	(0.05)	(0.05)	(0.05)
Grandmother	3.71	4.30	3.16	3.38
	(0.03)	(0.04)	(0.04)	(0.04)
Grandfather	4.89	5.47	4.33	4.54
	(0.03)	(0.05)	(0.04)	(0.04)
Maternal Height (cm)	151.1	151.5	150.9	151.0
	(0.04)	(0.06)	(0.06)	(0.06)
Paternal Height (cm)	162.0	162.6	161.6	162.0
	(0.05)	(0.07)	(0.06)	(0.06)

(continued on next page)

Table 1 *(continued)*

	Sample			
	All Children (Birth – 16 yrs)	Height-for-Age (Birth – 6 yrs)	Cognitive Score (7–16 yrs)	Attended Kindergarten (6–14 yrs)
Urban Household	0.51	0.53	0.49	0.50
	(0.00)	(0.01)	(0.01)	(0.01)
N. Households	8371	5392	5377	5281
N. Families	5263	4252	4163	3682
N. Children (Observations)	14558	6567	7727	7493

[*] In Rp0,000 (~1 USD).

Earnings are one potential candidate for individual-specific resources under the control of each family member. However, earnings reflect decisions about time allocation which is likely to be the outcome of a within-family bargaining process. Moreover, as a result of that process, income is typically shared with other household members and possibly also non co-resident family members through inter-household transfers.

Rather than rely on earnings, we assume that the distribution of assets within the family are the outcome of a less proximate bargaining process and, from the perspective of investments in young children, can be treated as pre-determined in the models. Exploiting an unusual feature of IFLS designed to test these models, we draw on information collected from each respondent age 15 or older about the value of assets the respondent owns in each of ten asset groups. Three of the asset groups are relatively illiquid – owner-occupied houses, other houses, and land – whereas the other seven are relatively liquid – livestock; vehicles; household appliances; financial assets; gold and jewelry; household furniture; and other assets. When assets are jointly owned with others inside or outside the family, the respondent estimates the value of the share of the asset that he or she owns.[7]

From the perspective of bargaining models, it is not clear that more and less liquid assets will be treated the same. In the empirical models, we separate liquid from illiquid assets for three reasons. First, from a theoretical perspective, if markets are complete, the distinction between liquid and illiquid assets should not matter. However, tests based on labor supply document that markets are not complete in Indonesia (LaFave and Thomas, 2016). The key illiquid assets in our study are land and housing. Markets for both assets are extremely thin. Very few sales of either asset are reported in IFLS. Access to credit is limited and it is not common for either land or housing to be used as collateral for loans. Moreover, rental markets are thin: only 2% of all households report any income from renting out land or housing. The vast majority of land and housing wealth is kept in the family and transfers are governed primarily by norms surrounding marriage and inheritance.

Second, from an empirical perspective, valuing illiquid assets is very difficult precisely because markets are thin. In contrast, liquid assets, such as cash, stocks and gold are very easy to value. Gold is of particular interest. It is stored in the form of simple jewelry, its price is set in world terms and, according to IFLS, there is a market for gold within 20 minutes of every village in the survey.

Third, our tests of resource allocation rely on attribution of asset values to individuals within families. Land and housing are typically reported as being jointly owned by a couple and jointly-owned assets are reported as being split equally among owners in over 90% of cases. In contrast, gold, stocks and cash are reported as primarily owned by individuals. For example, women own 80% of all gold which accounts for the vast majority of assets that are individually-owned by women. Furthermore, if a family member threatens to not co-operate, it may be difficult to withdraw his or her share of the value of the house from decision-making without evicting family members. Whether more and less liquid assets will have different impacts on child outcomes is an empirical question that will be explored below using the value of assets owned by each respondent as the measure of resources, y_m, in the models.

While individual-specific assets have considerable appeal in these models since individuals are likely to take the assets they own with them if they sever ties with the family, the assumption that the distribution of assets within the family is unrelated to unobserved factors that affect child outcomes is not without controversy.

The Indonesian context is important in this regard. Norms surrounding intra-family transfers of assets in Indonesia are governed by a combination of adat (community norms) and Islamic traditions. While there is heterogeneity in adat across the archipelago, according to the anthropological literature, the vast majority of assets are transferred across generations at the death of a parent, marriage of a child or birth of a child.

Land and housing are typically transferred to surviving children on the death of the parent. Given the thinness of these markets, and the relatively small plots, particularly in Java, it is very uncommon for parents to sell parts of plots of land while alive or give part of a plot to one child. The transfer of assets at marriage is widespread in Indonesia and it is typically in the form of gold and cash which is given to each individual getting married. In IFLS, 85% of all women report owning some assets at the time of marriage as do 75% of all men; for most of these people, those assets were given to them at the time of the marriage apart from a small amount of gold that is often given to a female at birth. The average value of assets brought to marriage is not significantly different for males or females (with males reporting 15% more assets). Importantly, it is normal for each individual to retain control over assets brought to the marriage, particularly among the Javanese, the dominant ethnic group (Bowen, 2003; Geertz, 1961; Contreras et al., 2002; Williams and Guest, 2012).

Evidence in IFLS is also consistent with the anthropological evidence that most intra-family wealth transfers in Indonesia occur around key demographic events and that large-scale, inter-generational, inter-vivos transfers at other times are not common. Specifically, IFLS respondents report the incidence and values of monetary and in-kind transfers to and from non-coresident parents, children and siblings over the previous 12 months. Transfers are common: 88% of households report sending transfers to other family members and 83% report receiving a transfer with 76% reporting both sending and receiving a transfer. This indicates there are extensive financial assistance links among non co-resident family members. However, the value of transfers is small relative to wealth. The median transfer sent to another family members if Rp 60,000 and the median received is Rp 16,000; since many households report both sending and receiving, it is of interest to examine the net transfer (sent less received) which is Rp 40,000 at the median and equal to 0.5% of median liquid assets (which are Rp 7.4 million). The incidence and value of net transfers are not related to household wealth, the wealth of individuals within a household, or the number of children in a household.

Turning to the empirical models of child human capital outcomes, the assumption that wealth is exogenous in the models raises three issues. First, there may be common unobserved factors that drive total family wealth and child human capital outcomes. This is a problem in much of the literature on child outcomes. A key strength of our approach is that we estimate models of child outcomes that include family fixed effects which absorb all common unobserved factors to the extent they affect child outcomes in a linear and additive way.

[7] Ninety-four percent of respondents provide their own estimate of the value of their assets. For the other 6%, the value reported by the spouse or household head is used.

Table 2
Empirical results.

Panel A: Model Estimates	*Family Fixed Effects*			*Household and Family Resources*			*Household and Family Resources*		
	Liquid and Illiquid Assets			*Liquid and Illiquid Assets*			*Liquid Assets*		
	Height- for-Age	**Cognitive Score**	**Attended Kindergarten**	**Height- for-Age**	**Cognitive Score**	**Attended Kindergarten**	**Height- for-Age**	**Cognitive Score**	**Attended Kindergarten**
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**	**(9)**
Liquid Household Assets	6.67	0.90	3.87	8.35	1.12	4.24	9.08	1.17	5.02
	(3.06)	(0.42)	(0.98)	(1.61)	(0.20)	(0.49)	(1.51)	(0.18)	(0.45)
Liquid Family Assets				2.29	0.37	1.45	3.43	0.45	1.94
				(1.73)	(0.23)	(0.56)	(1.61)	(0.22)	(0.52)
Illiquid Household Assets	1.59	−0.40	0.90	3.42	0.27	2.38			
	(3.22)	(0.44)	(1.01)	(1.80)	(0.23)	(0.55)			
Illiquid Family Assets				1.86	0.27	0.56			
				(1.27)	(0.16)	(0.46)			
Family Fixed Effects	Y	Y	Y	N	N	N	N	N	N
N. Observations	6567	7727	7493	6567	7727	7493	6567	7727	7493
Panel B									
B1. Combined Effects									
Household liquid plus illiquid	8.25	0.50	4.77	11.77	1.39	6.63			
coefficient (std. error)	(3.62)	(0.55)	(1.26)	(1.98)	(0.25)	(0.59)			
B2. Equality of Asset Effects (*p*-values)									
Household liquid and illiquid	0.32	0.05	0.05	0.08	0.01	0.03			
Liquid Assets				0.02	0.02	0.00	0.02	0.02	0.00
Illiquid Assets				0.50	0.99	0.01			
All Liquid and Illiquid				0.01	0.01	0.00			
B3. Coefficient Ratios: Household to family [...]									
Liquid Assets				3.64	3.03	2.92	2.65	2.60	2.60
				(2.93)	(2.03)	(1.23)	(1.40)	(1.40)	(0.77)
Illiquid Assets				1.84	1.01	4.24			
				(1.69)	(1.06)	(3.64)			

Panel C: Collective Model Nonlinear Wald Tests (*p*-values)			*p*-values		*p*-values
	Test of equality of ratios between [...] and [...]		Liquid Assets	Illiquid Assets	Liquid Assets
	Height-for-Age	Cognitive Score	0.86	0.66	0.98
	Height-for-Age	Attend Kindergarten	0.80	0.50	0.97
	Cognitive Score	Attend Kindergarten	0.96	0.26	0.99
		All Ratios	0.99	0.68	0.99

Notes: Panel A: All regressions include assets in log form, as well as controls for gender, age, household and family size and composition, parental education, age and gender of the household and family head and location as described in the text. Standard errors in parenthesis account for clustering at the family level. Panel B1 reports the sum of household liquid and illiquid assets effects and its standard error; panel B2 reports *p*-values for tests of the equality of asset effects in each row; panel B3 reports ratios and their standard errors calculated using the delta method. Panel C: The *p*-values correspond to results from nonlinear Wald tests where the null is Pareto efficiency. These tests compare the equivalence of the ratios shown in Panel B3 for the outcome pairs in each row.

Second, the distribution of wealth within the family may be endogenous with respect to child human capital outcomes. However, the mechanisms underlying unobserved factors that affect child outcomes and the distribution of assets within the family are not clear. There are multiple mechanisms that might underlie decisions about asset transfers. To illustrate with an example, a grandmother may give assets to a son or daughter to assure that child invests more in the grandchild, or the grandmother may hold back assets and invest in the grandchild herself, rather than trust the child. It is also possible that a grandmother provides assets as a reward for a grandchild having better human capital outcomes, or the grandmother may be inequality-averse and provide assets to compensate for poor outcomes of a grandchild. In all of these cases, it is not clear that the grandmother is better off giving assets to a child unless she and the child share the same preferences. Put another way, in the unitary model, the distribution of wealth is irrelevant. In an efficient model with preference heterogeneity, assets may be distributed to assure that all allocations are efficient – in which case the results of the model will reflect this fact – or assets may be used to assure spending on goods and services achieves that end.

Whatever the mechanism, it is possible that the distribution of wealth within a family will be correlated with the presence of a child in a specific household and that, in turn, the distribution of family wealth between that household and other households will be correlated with unobserved characteristics that affect child human capital outcomes. To address this concern, we estimate the human capital function with family fixed effects, (7) and compare the impact on child outcomes of assets of members of the child's household with assets of other non co-resident family members. To the extent that the key source of endogeneity lies in the distribution of assets between households, these estimates should not be contaminated by unobserved heterogeneity. Moreover, household-level assets may be measured with less error and so there will be gains in precision relative to estimating model (5) with individual-specific resources. This approach is most comparable with Altonji et al. (1992) and Hayashi et al. (1996) and others in this literature who have not examined the impact of individual-level family resources but, instead, examined the impact of household expenditures relative to the impact of expenditures of other family members who are not co-resident. That approach does not have a strong link to theory as the distribution of expenditure between households in the family is likley to be endogenous if there is sharing of resources.

As an alternative approach to addressing endogeneity of the distribution of wealth in the family, we note that the anthropological literature highlights the key role of assets at marriage in determining the resources over which a person has control. It is reasonable to treat the distribution of those assets as exogenous with respect to observed child human capital outcomes (since less than 0.2% of all births are premarital in Indonesia). This suggests that the real value of assets brought to the marriage are potentially valid instruments for the distribution of wealth within the family at the time of the survey. IFLS was designed with this in mind and asked every respondent who had ever been married the value of their own assets at the time of the marriage and the value of the dowry paid. Results of experiments using assets at marriage as instruments for liquid assets, described below, indicate that exogeneity of assets in models of human capital outcomes cannot be rejected. This is, perhaps, not surprising given that inter-vivos wealth transfers are not common.

The distribution of assets within families is reported, along with other summary statistics in Table 1. Since the samples vary across outcomes (because age-eligibility applies to each outcome), the statistics are reported for each sample. Overall, the average family has Rp206 million combining illiquid and liquid assets (which is about $20,000) with the average household in which a target child lives accounting for about forty percent of the total. Wealth is split between liquid holdings, which account for approximately twenty-two percent of total assets, and illiquid holdings. Within couples, husbands tend to report slightly higher levels of liquid assets than wives. Child characteristics, X_i, and characteristics of other family members, X_m, are reported in the rest of Table 1 and serve as controls in all regression models.[8]

4. Empirical results

Regression results of the effects of resources on child outcomes are reported in Table 2 for the models with resources measured at the household level using the logarithm of household and other family assets (in Rp 10,000s).[9] Results for models with individual-specific resources are reported in Table 3.

To begin, columns 1 through 3 of Table 2 present results of estimating (9), the model with a family fixed effect that measures assets at the household level for each of the three child outcomes. The fixed effect absorbs family-level characteristics that have a linear and additive impact on child outcomes. As this includes total family resources, an increase in household resources can be interpreted as an increase in the share of total family resources attributed to the household. If the unitary model of family decision-making is correct, the distribution of resources within the family should have no impact on child outcomes and so the resources of the household in which the child resides should not be significantly related to any of the child outcomes.

The empirical model distinguishes assets that are more liquid from those that are less liquid (or illiquid) and includes the logarithm of the value of household assets in each class. Liquid assets have a statistically significant positive impact on each of the child outcomes at a 5% size of test; in contrast, illiquid assets are not predictive of child outcomes. As shown in panel B1, taken together, liquid and illiquid assets are significantly associated with height for age and kindergarten attendance. If illiquid assets are not included in the model, then the estimated coefficients and standard errors for liquid assets are essentially the same. The evidence in these models is not consistent with the unitary model of the family.

A key advantage of the model with a family fixed effect is that it is not necessary to measure total family resources. Columns 4 through 6 of Table 2 present results from estimating model (7) which includes both the logarithm of household assets, y_{hf}, and the logarithm of assets of family members who are not co-resident with the target child, $y_{\bar{h}f}$, again distinguishing liquid from illiquid assets. Holding family resources constant, the logarithm of liquid household assets has a significant positive impact on each of the three child outcomes but illiquid household assets affect only kindergarten attendance. Taking household liquid and illiquid assets together, household resources are significant predictors of child outcomes (as shown in panel B1 of the table). In all cases, the impact of family resources is smaller and not significant in all but the case of illiquid assets which affect kindergarten attendance.

If family behavior is unitary, the effect on child outcomes of household resources and resources of non co-resident family members will be the same. Tests of equality of the estimated effects are shown in panel B2. Whereas there are no significant differences in the impacts of

[8] Controls include gender and age of the child, using indicator variables for each age, and whether the child lives with their mother and father. At the household, and family level when appropriate, we include demographic controls including household size and composition, and age and education of the household and family head. Factors common at the province level are controlled for with province fixed effects, and we include an indicator for whether the household is located in an urban or rural region. In regressions examining child height, both mother's and father's height is included to capture the genetic component of height as well as additional factors not controlled for by parental education and age.

[9] Only 0.2% of households report no assets. Since the logarithmic transformation of assets is very convenient for interpretation of the coefficient estimates, households that report no assets are assigned the sample mean and identified by an indicator variable in the regression models. None of the conclusions depends on this choice; models that use the quartic root of assets and hyperbolic inverse sine transformation yield the same inferences.

illiquid assets of households and the rest of the family for height-for-age and cognitive scores, for liquid assets, and also for all assets taken together, the unitary model is rejected at a 5% size of test for each of the child outcomes.

It is not obvious how a family boundary should be delineated. As a practical matter, we have defined families as related biological kin of the parents of the target child. However, our empirical measure of family resources includes only the resources of the kin of one of those parents, the parent who was a member of the baseline household and whose parents and siblings have been interviewed. For all but a handful of children, the other parent is a new entrant into the IFLS sample and other members of that person's family are not interviewed. This is a function of the fact that the data are collected as part of a household-based rather than family-based survey. Measuring resources for only one side of the target child's family should not bias our results since whether resources are from the maternal or paternal side is essentially random. Moreover, including a control for whether we observe the maternal or paternal side of a child's family tree does not affect the results, nor do the asset effects statistically differ by the interaction of gender of the child and the maternal or paternal line.

A related concern is that family resources are measured only for those kin who are members of households that are included in IFLS. One measure of the importance of this concern is provided by comparing the estimated difference between the impact of household assets in the models that control family assets (in columns 4 through 6) with the impact of household assets in the models that include a family fixed effect (in columns 1 through 3) and so are not contaminated by measurement error in family assets. For child height for age, the difference between the impact of the logarithm of household and family liquid assets is 6.06 and the fixed effects estimate of this difference is 6.67; for illiquid assets the difference in the model that includes measures of family resources is 1.56 and in the fixed effects model it is 1.59. Neither of these differences is statistically significant nor economically meaningful. The estimates for the cognitive score and kindergarten attendance are also not statistically significant. We conclude that measurement error in family resources is not an empirically important issue in these models.

The impacts of household and family assets that are illiquid are small in magnitude relative to liquid assets and statistically significant in only one case indicating that, conditional on liquid assets, variation in illiquid assets is not an important determinant of child outcomes. Individuals and households who have more liquid assets tend to have more illiquid assets which likely affects the precision of the estimated effects of resources. The models are, therefore, re-estimated dropping illiquid assets. The results are reported in columns 7 through 9. The coefficient estimates are slightly larger and better determined than the models with illiquid assets. All of the estimated effects of household and family liquid assets are statistically significant and household assets have a significantly bigger effect on each child outcome than family assets indicating unambiguous rejection of the unitary model.

The ratios of the effects of own household resources to resources of non co-resident family members, (8), are presented in columns 7 through 9 of panel B3 of the table along with *p*-values of non-linear Wald tests of the equality of those ratios in panel C. The ratios for each child outcome are all very close, precisely estimated, and indicate that resources of those who co-reside with the child have about 2.6 times the impact relative to resources of non co-resident family members. Given the similarity of all three ratios, it is not surprising that none are different for any pair of outcomes or for all three taken together as shown by the test statistics in column 9 of panel C. The similarlity of the ratio estimates aleviates concerns of weak power in the non-linear Wald test as well. Whereas the unitary model of the family has been rejected by these data, we cannot reject the collective model of family behavior: with respect to child investments, co-resident and non co-resident family members appear to co-ordinate resource allocations Pareto efficiently. The same conclusion is drawn for the models with liquid and illiquid assets although in those cases the ratios are not as well determined.[10]

An advantage of a model of the family rather than the household to better understand resource allocations is that the family model does not need to take into account either potentially endogenous living arrangements or make strong assumptions about selection of household members. We turn, therefore, to the family-based model (5) and report estimates of the effects of resources of parents and grandparents on the three child outcomes in Table 3. Since the power of the Pareto efficiency tests is a function of the precision of the estimated resource effects, we include only liquid assets in these models.

Maternal resources have the largest impact on the child outcomes and the estimated effects are statistically significant in all three cases, a result that is consistent with a large literature (Thomas, 1990). The effects of paternal resources are also positive and while they are only significant in the case of kindergarten attendance, taken together maternal and paternal resources are significant predictors of each child outcome (with a p-value < 0.01 in each case). The effects of resources of grandparents are smaller in magnitude, none are statistically significant individually or taken together. The effects of grandparents resources are significantly different from maternal resources for all three child outcomes. The resources of other household members are significant predictors of each child outcome, the estimated effects are larger (albeit insignificantly) than paternal resources for child height but significantly smaller than the effects of maternal resources on cognition and kindergarten attendance. The same pattern is apparent in the final row of panel A of the table which reports the effects of resources of other, non co-resident family members. There are no substantively important or statistically significant differences between the effects of other family members who do or do not co-reside with the child.[11] These family members are aunts and uncles of the target child and, in some cases, older siblings.

The evidence in Table 3 demonstrates that, in terms of allocating resources for child human capital, the unitary model is rejected not only for the family but also for the household. While those are not novel conclusions, the results in Table 3 also demonstrate that the relationships between resources of different family members and child outcomes is extremely complex, not driven by living arrangements, and points to the likely value-added of designing studies that reach beyond parents and household members to more fully understand family dynamics, at least in a low income setting like Indonesia.

With six demographic groups, there are 15 resource effect ratios for each outcome. Given the centrality of maternal resources in the estimates reported in panel A of the table, we report the ratios relative to the effect of maternal assets for each child outcome as an illustration in panel B. As the impact of maternal resources is the largest for each outcome, all of the ratios are less than one. Following (5), if family decisions are Pareto efficient, these ratios will not be different from each other. Appendix Table A2 reports p-values for the equality of each pair of ratios and the p-values for the equality of all the ratios for each child outcome pair are repeated in panel C of Table 3. None of the ratios is significantly different from another in pairwise tests, nor tested

[10] As noted above, unobserved factors that drive the distribution of assets within the family may be correlated with unobserved factors that drive child human capital outcomes. We have investigated the importance of this concern using the real value of assets brought to marriage as instruments for current assets in the models in the table. Assets at marriage are strong predictors of household and extended family liquid assets at the time of the survey: the F statistics range from 59.9 to 72.4. Over-identification tests fail to reject the validity of the instruments. However, the IV and OLS estimates are not significantly different from each other. Specifically, we fail to reject the hypothesis that all the coefficients in the IV and OLS models are the same with p-values of 0.99; focusing only on the coefficients on assets, we also fail to reject equality of the IV and OLS estimates with p-values of 0.10 for height-for-age, 0.51 for cognitive scores and 0.11 for kindergarten attendance. We choose, therefore, to focus on the OLS results.

[11] In models that include both illiquid and liquid assets, we find that illiquid assets of the father and grandparents are significant predictors of kindergarten attendance but neither height nor cognition. This may reflect the fact that kindergarten attendance is readily observed and its timing is known at child birth.

Table 3
Individual level results.

Panel A: Model Estimates	*Individual Resources - Liquid Assets*		
	Height-for-Age	**Cognitive Score**	**Attended Kindergarten**
	(1)	**(2)**	**(3)**
Mother's Assets	4.65	0.82	2.76
	(1.81)	(0.21)	(0.58)
Father's Assets	2.81	0.31	2.12
	(1.79)	(0.21)	(0.57)
Grandmother's Assets	2.18	0.10	1.10
	(1.90)	(0.31)	(0.75)
Grandfather's Assets	0.20	0.05	−0.23
	(2.00)	(0.34)	(0.87)
Rest of Household Assets	3.28	0.29	0.99
	(1.50)	(0.16)	(0.42)
Rest of Family Assets	3.06	0.38	1.99
	(1.70)	(0.21)	(0.51)
N. Observations	6567	7727	7493
Panel B: Coefficient Ratios relative to Mothers			
Mother	1.00	1.00	1.00
	–	–	–
Father	0.60	0.38	0.77
	(0.56)	(0.33)	(0.33)
Grandmother	0.47	0.12	0.40
	(0.46)	(0.38)	(0.28)
Grandfather	0.04	0.06	−0.08
	(0.43)	(0.42)	(0.32)
Rest of Household	0.71	0.36	0.36
	(0.45)	(0.21)	(0.17)
Rest of Family	0.66	0.47	0.72
	(0.46)	(0.29)	(0.25)

Panel C: Collective Model Nonlinear Wald Tests (*p*-values)

Equality of all ratios between [...] and [...]		*p*-value
Height-for-Age	Cognitive Score	0.99
Height-for-Age	Attend Kindergarten	0.99
Cognitive Score	Attend Kindergarten	0.99
	All Ratios	0.96

Notes: Each regression includes log liquid assets at the individual level and controls for the age and gender of the child, household and family size and composition, age and education of parents and grandparents, and location as described in the text. Standard errors in parentheses account for clustering at the family level. Panel B reports coefficient ratios and their standard errors with the effect of mother's liquid assets serving as the denominator. Panel C reports *p*-values from nonlinear Wald tests across all ratio combinations of the outcome pairs in each row.

jointly.

Taken together, the evidence in Tables 2 and 3 indicates that it is not possible to reject the hypothesis that family behavior is consistent with Pareto efficiency. Apparently, non co-resident family members successfully navigate information asymmetries, market imperfections and barriers to exchange so that, at least with regard to child human capital investments, resources are allocated efficiently.

Sensitivity analysis and robustness. In order to ensure the validity of the results, numerous robustness checks were conducted to examine the failure to reject efficiency and rejection of the unitary model in a number of subsamples and alternative specifications to address various hypothesis concerning the sharing habits of extended families. Our conclusion that family behavior is consistent with the collective model is upheld when we allow nonlinearities in the impact of resources, consider other child outcomes or examine specific subsamples of children.

We have tested whether the effects of resources on child outcomes are non-linear by estimating models with splines and models with polynomials; the linearity assumption is not rejected for each outcome. Moreover, our conclusion that families are co-operative is not changed in these non-linear models. The conclusions also extend to other child outcomes including weight for age, body mass index and school enrolment rather than early attendance.

It is possible that some sub-populations do not behave efficiently. For example, families with more wealth may be able to afford to deviate from efficiency, or co-ordination failures may be more likely in larger or more dispersed families. To explore whether Pareto efficiency describes the behavior of these population sub-groups, we stratified the sample based on total family wealth, on family size, the number of children within a family and whether families were geographically dispersed. In none of these sub-samples was Pareto efficiency rejected.

While the IFLS is an extremely rich resource for this research, it and every other household-based data collection effort is limited in that resources of all extended family members are not recorded. We established that our estimates in models that control family resources and models that treat family resources as an unobserved characteristics are very similar and not statistically significantly different. This gives us some comfort. We have also exploited the fact that, by design, there is variation in the branch of the family tree included in the IFLS: for some children, the mother's family was selected for the baseline and so the

resources of her family are measured; for other children, it is the paternal side of the family that has been interviewed. To assess whether it matters if the root household is on the maternal or paternal side of the family, we separately estimated models for families where the mother in a new household is an IFLS panel member compared with those where the father is an IFLS panel member. Again, we found no significant differences. In addition, IFLS collects demographic information about all parents and siblings of each respondent; we can control for whether the entire (paternal or maternal) family has been enumerated in IFLS and we can restrict our sample to those families in which the majority of family members have been interviewed. Our conclusions are not changed.

5. Conclusion

The family plays a central role in many models of behavior in economics, and both implicit and explicit transfers between non co-resident family members have been shown to play an important role in many societies, particularly in lower income settings. Whereas understanding of household behavior has been substantially enriched by recent advances in the theoretical and empirical literatures on decision-making at the household level, relatively little is known about the behavior of extended family members who are not co-resident. Extending theoretical models of household decision-making to the family context, we draw on extremely rich survey data from Indonesia that was designed to collect individual-level measures of economic resources of extended family members in order to empirically discriminate among different models of family behavior.

We examine the relationship between the distribution of assets within a family and three child outcomes that are related to human capital: height for age, cognition and kindergarten attendance. Not only are parents likely to care about the human capital of their children, but other family members are also likely to value higher levels of human capital among the next generation. Indeed, we find parental resources – particularly maternal resources – as well as resources of adults who co-reside with the target child have a positive impact on human capital outcomes. We also find these investments are greater if, conditional on household resources, non co-resident family members have more assets, although these effects are significantly smaller in magnitude than the effects of parental and household assets. This behavior is not consistent with family members pooling all resources and behaving as a single unit. However, we cannot rule out that extended families behave cooperatively and that their allocation decisions are Pareto efficient, at least with respect to investments in children.

The results are important for at least three reasons. First, they establish that resources of the extended family play a role in early life human capital accumulation over and above the resources of a child's household. Second, in a low income context where market imperfections and liquidity constraints are likely to be important, geographically dispersed families behave as it resource allocations are coordinated (Pareto) efficiently in spite of informational asymmetries and transaction costs. Third, taken together, these results suggest that there are likely to be substantial benefits to moving beyond surveys of households and developing studies that systematically collect data on family members in order to better understand the roles that families actually play in societies today.

Appendix A

See Tables A1 and A2 here

Notes: Table reports *p*-values from nonlinear Wald tests of Pareto efficiency. Each cell is the *p*-value comparing the ratio of marginal effects for assets owned by the individual listed in the row with the individual in the column. For example, 0.81 is the *p*-value for the test of equivalence between the ratio of effects of assets controlled by a child's mother to their father for height-for-age and cognitive scores. Also included in each

Table A1
Sample description.

Number of unique [...]	
Children *(birth – 16 yrs)*	14881
Households	8367
Families	5283
Mothers	8438
Fathers	7809
Grandmothers	5448
Grandfathers	4386

Table A2
Individual level tests of the collective model (*p*-values).

	Outcome Pair				
	Father	Grandmother	Grandfather	Rest of HH	Rest of Fam
			Height-for-Age and Cognitive Score		
Mother	0.72	0.56	0.97	0.45	0.71
Father		0.74	0.95	0.86	0.93
Grandmother			0.91	0.80	0.70
Grandfather				0.93	0.95
Rest of Household					0.76
Equality among parents and grandparents (6 ratios)			0.99		
Equality among all ratios (15 ratios)					0.99

(continued on next page)

Table A2 *(continued)*

	Outcome Pair				
	Father	Grandmother	Grandfather	Rest of HH	Rest of Fam
			Height-for-Age and Attended Kindergarten		
Mother	0.81	0.89	0.81	0.41	0.91
Father		0.77	0.83	0.33	0.87
Grandmother			0.80	0.70	0.85
Grandfather				0.76	0.82
Rest of Household					0.38
Equality among parents and grandparents (6 ratios)			0.99		
Equality among all ratios (15 ratios)					0.99
			Cognitive Score and Attended Kindergarten		
Mother	0.41	0.55	0.76	0.99	0.51
Father		0.86	0.81	0.49	0.78
Grandmother			0.82	0.57	0.76
Grandfather				0.77	0.79
Rest of Household					0.63
Equality among parents and grandparents (6 ratios)			0.97		
Equality among all ratios (15 ratios)					0.99
Overarching Test					
Equality across all ratios, all outcomes (45 ratios)					0.96

panel are tests of equivalence of ratios for parents and grandparents, and the test of equality of all fifteen ratios in the panel. Rest of HH and Rest of Fam stand for Rest of Household and Family. The overarching test at the bottom of the table is for equality of all ratios across all outcomes (45 total ratios). Tests are conducted allowing for clustering at the family level.

References

Alderman, H., Hoddinott, J., Kinsey, B., 2006. Long term consequences of early childhood malnutrition. Oxford Econ. Pap. 58 (3), 450–474.

Altonji, J., Hayashi, F., Kotlikoff, L., 1992. Is the extended family altruistically linked? Direct tests using microdata. Am. Econ. Rev. 82 (5), 1177–1198.

Angelucci, M., de Giorgi, G., 2009. Indirect effects of an aid program: how do cash transfers affect ineligibles' consumption? Am. Econ. Rev. 99 (1), 486–508.

Angelucci, M.G., de Giorgi, M., Rangel, Rasul, I., 2010. Family networks and school enrollment: evidence from a randomized social experiment. J. Public Econ. 94 (3–4), 197–221.

Antman, F., 2011. The intergenerational effects of paternal migration on schooling and work: what can we learn from children's time allocations? J. Dev. Econ. 96 (2), 200–208.

Ashraf, N., 2009. Spousal control and intra-household decision-making: an experimental study in the Philippines. Am. Econ. Rev. 99 (4), 1245–1277.

Becker, G., 1981. A Treatise on the Family. Harvard University Press, Cambridge, Massachusetts.

Bobonis, G., 2009. Is the allocation of resources within the household efficient? New evidence from a randomized experiment. J. Political Econ. 117 (3), 453–503.

Bourguignon, F., Browning, M., Chiappori, P.A., Lechene, V., 1993. Intra household alocation of consumption: a model with some evidence from French data. Ann. D'Econ. et De. Stat. 29, 137–156.

Bowen, J., 2003. Islam, Law and Equality in Indonesia: An Anthropology of Public Reasoning. Cambridge University Press, New York.

Browning, M., Chiappori, P.A., 1998. Efficient intra-household allocations: a general characterization and empirical tests. Econometrica 66 (6), 1241–1278.

Browning, M., Bourguignon, F., Chiappori, P.A., Lechene, V., 1994. Incomes and outcomes: a structural model of intra-household allocation. J. Political Econ. 102 (6), 1067–1096.

Cameron, L., Cobb-Clark, D., 2008. Do coresidency and financial transfers from the children reduce the need for elderly parents to works in developing countries? J. Popul. Econ. 21 (4), 1007–1033.

Chiappori, P.A., 1988. Rational household labor supply. Econometrica 56 (1), 63–89.

Chiappori, P.A., 1992. Collective labor supply and welfare. J. Polit. Econ. 100 (3), 437–467.

Chiappori, P.A., 1997. Introducing household production in collective models of labor supply. J. Polit. Econ. 105, 191–209.

Chiappori, P.A., Mazzocco, M., 2016. Static and intertemporal household decisions. J. Econ. Lit., (Forthcoming).

Contreras, D., Thomas, D., Frankenberg, E., 2002. Distribution of power within the household and child health. Mimeo.

Cox, D., Jimenez, E., 1990. Achieving social objectives through private transfers: a review. World Bank Res. Obs. 5 (2), 205–218.

Dalton, M., Hotz, J.V., Thomas, D., 2016. Resources, composition, and family-decision making. Mimeo.

Duflo, E., 2003. Grandmother and granddaughters: old-age pensions and intrahousehold allocations in South Africa. World Bank Econ. Rev. 17 (1), 1–25.

Fafchamps, M., Lund, S., 2003. Risk-sharing networks in rural Philippines. J. Dev. Econ. 71 (2), 261–288.

Farfan, G., Genoni, M., Thomas, D. , Velasquez, A., 2016. Extended Families in Mexico and the United States, Mimeo.

Fogel, R., 2004. The Escape From Hunger and Premature Death, 1700–2010: Europe, America and the Third World. University of Chicago Press, Chicago.

Frankenberg, E., Karoly, L., 1995. The 1993 Indonesia family life survey: overview and field report (Technical Report). Rand Corp..

Frankenberg, E., Thomas, D., 2000. The Indonesia Family Life Survey (IFLS): Study Design and Results from Waves 1 and 2 (Technical Report). Rand Corporation, Santa Monica, CA.

Frankenberg, E., Lillard, L., Willis, R., 2002. Patterns of intergenerational transfers in Southeast Asia. J. Marriage Fam. 64 (3), 627–641.

Frankenberg, E., Smith, J.P., Thomas, D., 2003. Economic shocks, wealth, and welfare. J. Hum. Resour. 38 (2), 280–321.

Geertz, H., 1961. The Javanese Family: A Study of Kinship and Socialization. Free Press, Glencoe, IL.

Geertz, H., 1963. Indonesian Cultures and Communities. Hraf Press, New Haven, CT.

Genoni, M., 2012. Health shocks and consumption smoothing: evidence from Indonesia. Econ. Dev. Cult. Change 60, 3.

Gertler, P., Gruber, J., 2002. Insuring consumption against illness. Am. Econ. Rev. 92 (1), 51–70.

Gertler, P., Heckman, J., Pinto, R., Zanolinim, A., Vermeersh, C., Walker, S., Chang, S., Grantham-McGregor, S., 2014. Labor market returns to an early childhood stimulation intervention in Jamaica. Science 344 (6187), 998–1001.

Hamoudi, A., Thomas, D., 2014. Endogenous coresidence and program incidence: South Africa's old age pension. J. Dev. Econ. 109, 30–37.

Hayashi, F., Altonji, J., Kotlikoff, L., 1996. Risk-sharing between and within families. Econometrica 64 (2), 261–294.

Heckman, J., 2006. Skill formation and the economics of investing in disadvantaged children. Science 312 (5782), 1900–1902.

Heckman, J., Kautz, T., 2012. Hard evidence on soft skills. Labour Econ. 19 (4), 451–464.

Hoddinott, J., Behrman, J., Maluccio, J., Meglar, P., Quisumbing, A., Ramirez-Zea, M., Stein, A., Yount, K., Martorell, R., 2013. Adult consequences of growth failure in early childhood. Am. J. Clin. Nutr. 98 (5), 1170–1178.

Jay, R., 1969. Javanese Villager: Social Relations in Rural Madjokuto. MIT Press, Cambridge, MA.

Jensen, R., 2003. Do private transfers 'displace' the benefits of public transfers? Evidence from South Africa. J. Public Econ. 88 (1–2), 89–112.
Kaplan, R., Saccuzzo, D., 1997. Psychological Testing: Principles, Applications, and Issues 4th ed.. Brooks/Cole, Pacific Grove, CA.
Knudsen, E.I., Heckman, J.J., Cameron, J.L., Shonkoff, J.P., 2006. Economic, neurobiological and behavioral perspectives on building America's future workforce. Proc. Natl. Acad. Sci. 103, 10155–10162.
Kreager, P., Schroder-Butterfill, E., 2007. Gaps in the family networks of older people in three Indonesian communities. J. Cross-Cult. Gerontol. 22 (1), 1–25.
LaFave, D., Thomas, D., 2016. Farms families and markets: new evidence on the completeness of markets in agricultural settings. Econometrica 85 (4), 1917–1960.
Lundberg, S., Pollak, R., 1996. Bargaining and distribution in marriage. J. Econ. Perspect. 10 (4), 39–58.
Lundberg, S., Pollak, R., Wales, T., 1997. Do husbands and wives pool their resources? Evidence from the United Kingdom child benefit. J. Hum. Resour. 32 (3), 463–480.
Martorell, R., Habicht, J., Growth in early childhood in developing countries, In: Falkner, F., Tanner, J., (eds.), Human Growth: A Comprehensive Treatise, Vol. 3, 2nd ed., Plenum Press, New York.
Munshi, K., 2003. Networks in the modern economy: Mexican migrants in the US labor market. Q. J. Econ. 118 (2), 549–599.
Nobles, J., 2013. Migration and father absence: shifting family structure in Mexico. Demography 50 (4), 1303–1314.
Okten, C., Osili, U.O., 2004. Social networks and credit access in Indonesia. World Dev. 32 (7), 1225–1246.
Park, C., 2003. Interhousehold transfers between relatives in Indonesia: determinants and motives. Econ. Dev. Cult. Change 51 (4), 929–944.
Rangel, M., 2006. Alimony rights and intrahousehold allocation of resources: evidence from Brazil. Econ. J. 116 (513), 627–658.
Raut, L.K., Tran, L.H., 2005. Parental human capital investment and old-age transfers from children: is it a loan contract or reciprocity for Indonesian families? J. Dev. Econ. 77 (2), 389–414.
Raven, J., 2000. The Raven's progressive matrices: change and stability over culture and time. Cogn. Psychol. 41, 1–48.
Rosenzweig, M., Stark, O., 1989. Consumption smoothing, migration and marriage: evidence from rural India. J. Polit. Econ. 97 (4), 905–926.
Rubalcava, L., Teruel, G., Thomas, D., 2009. Investment, time preferences and public transfers paid to women. Econ. Dev. Cult. Change 57 (3), 307–338.
Samuelson, P., 1956. Social indifference curves. Q. J. Econ. 70 (1), 1–22.
Schroder-Butterfill, E., 2005. The impact of kinship networks on old-age vulnerability in Indonesia. Ann. Demogr. Hist. 2.2, 139–164.
Silventoinen, K., Kaprio, J., Lahelma, E., Koskenvuo, M., 2000. Relative effect of genetic and environmental factors on body height: differences across birth cohorts among finnish men and women. Am. J. Public Health 90, 627–630.
Smith, J.P., Thomas, D., 1998. On the road: marriage and mobility in Malaysia. J. Hum. Resour. 33 (4), 805–832.
Strauss, J., Thomas, D., 1998. Health, nutrition, and economic development. J. Econ. Lit. 36 (2), (766, 817).
Strauss, J., Witoelar, F., Sikoki, B., Wattie, A.M., 2009. The Fourth Wave of the Indonesia Family Life Survey (IFLS4): Overview and Field Report (Technical Report). Rand Corporation, Santa Monica, CA.
Strauss, J., Beegle, K., Sikoki, B., Dwiyanto, A., Herawati, Y., Witoelar, F., 2004. The Third Wave of the Indonesia Family Life Survey (IFLS): Overview and Field Report (Technical Report). Rand Corporation, Santa Monica, CA.
Thomas, D., 1990. Intrahousehold resource allocation: an inferential approach. J. Hum. Resour. 25 (4), 635–664.
Thomas, D., Frankenberg, E., 2007. Household responses to the financial crisis in Indonesia: longitudinal evidence on poverty, resources and well-being. In: Ann Harrison (Ed.), Globalization and Poverty. University of Chicago Press, Chicago, IL.
Thomas, D., Witoelar, F., Frankenberg, E., Sikoki, B., Strauss, J., Sumantari, C., Suriastini, W., 2012. Cutting the costs of attrition: results from the Indonesia family life survey. J. Dev. Econ. 98, 108–123.
Townsend, R., 1994. Risk and insurance in village India. Econometrica 62 (3), 539–591.
Waterlow, J., Buzina, R., Keller, W., Lane, J., Nichman, M., Tanner, J., 1977. The presentation and use of height and weight data for comparing the nutritional status of groups of children under the age of ten years. Bull. World Health Organ. 55, 489–498.
Williams, Lindy, Guest, Michael Philip (Eds.), 2012. Demographic Change in Southeast Asia: Recent Histories and Future Directions. Cornell University Press, Ithaca.
Witoelar, F., 2013. Risk sharing within the extended family: evidence from the Indonesia family life survey. Econ. Dev. Cult. Change 62 (1), 65–94.