California Center for Population Research
University of California - Los Angeles

# Pension Income and the Well-Being of Children and Grandchildren: New Evidence from South Africa

**Amar Hamoudi**
**Duncan Thomas**

**CCPR-043-05**

**October 2005**

*California Center for Population Research*
*On-Line Working Paper Series*

# Pension Income and the Well-being of Children and Grandchildren: New Evidence from South Africa

Amar Hamoudi
(UCLA)

and

Duncan Thomas
(UCLA)

First version: March 2005
This version: October 2005

This research was supported by NICHD grant R01-HD040245. We are grateful to Statistics South Africa, the Medical Research Council of South Africa, Macro International, and the University of Natal for providing us with access to the data. We have benefited from discussions with Angus Deaton, Arie Kapteyn, and Dean Yang.

**Abstract**

*In the early 1990s, the South African Old Age Pension was expanded to cover most black South Africans above a gender-specific age cut-off. This expansion resulted in a substantial and arguably exogenous increase in the income of older South Africans. A series of very creative studies have exploited this source of variation in income to shed light on the ways in which familes and households allocate resources among their members. A key assumption underlying these studies is that pension income has no impact on unmeasured characteristics of those people who co-reside with pension recipients. This paper provides empirical evidence on the importance of this assumption. Pension-eligible adults are more likely to co-reside with other adults who have lower levels of human capital as measured by height and education. Since height and education are fixed for adults, this cannot be an effect of the pension income but rather reflects the selection of adults who co-reside with older adults when they become eligible for the pension. The paper proceeds to explore the importance of treating living arrangements as endogenous for re- interpretation of results on the impact of the pension in the literature. The evidence highlights the potential value of moving beyond theory and data which are bound by the confines of a spatially determined definition of the household.*

# Introduction

The social science literature is replete with attempts to measure the welfare impacts of some public program. In some cases, these evaluations are based on carefully designed randomized treatment-control designs. The majority, however, are conducted in non-experimental settings, requiring analysts to develop clever and innovative ways to isolate the causal relationship between the program and outcomes of interest. (See, for example, Moffitt, 1991, 2005; Heckman, Lalonde, and Smith, 1999, for discussions). The Old Age Pension (OAP) in South Africa is an extraordinarily generous publicly financed cash transfer program which was implemented in a non-experimental setting, but which nonetheless offers great promise in isolating its effects. It is non-contributory, and the means test is sufficiently generous that the

transfer is available to almost all of South Africa's majority black population when they reach the minimum age of eligibility (60 years for females, 65 for males). Exploiting these features of the program, several innovative studies have examined the impact of the OAP in South Africa on the well-being of individuals within recipient households. These studies not only have provided important insights into the impact of the program on welfare in South Africa, but also have contributed to the more general understanding of the ways in which families and households allocate resources among their members.

The approach in most of these studies involves taking advantage of the age eligibility rules. Specifically, since a person's date of birth is uncorrelated with any relevant individual characteristics, discontinuous differences in behaviors or outcomes between individuals or households which are just age eligible and those which are almost age eligible can be causally attributed to the OAP. Using this class of approaches, studies have found for example that in households with a woman of pension eligible age relative to those with a woman just short of pension eligibility, resident young children–and in particular young girls–are in better health as measured by anthropometric indicators (Deaton and Case, 1998; Duflo, 2000, 2003), that working-age men supply less market labor (Bertrand, Miller, and Mullainathan, 2004), and that private remittances into the household are lower (Jensen, 2004). This research points to two main conclusions–first, the public transfers have benefited individuals other than the primary recipient, especially when the primary recipient is a woman, and secondly that in South Africa, households do not behave as if they are "unitary" in the sense that the allocation of resources within households varies depending on the individual to whom they are attributed.

It is well understood, however, that the interpretation of these empirical findings depends on the nature of the relationship between household composition and the pension. This is likely to be particularly relevant in the South African context, where families have historically been distributed across multiple households, in part because of apartheid. Many rural families depend largely on transfers from family

members in urban areas for income, even married couples often live apart, and it is common for non co-resident kin to retain social and economic ties with origin households. Therefore, empirical evidence on the time allocation or consumption behavior of *households* is an incomplete reflection of the *family*'s allocation decisions. Indeed in one of the seminal papers in this literature, Case and Deaton highlighted this point, but noted that the data available at that time only covered a short span of time since the institution of the OAP in its current form, and so "the available data are unlikely to be informative about this question." As the body of available data has expanded, however, recent research has suggested that the composition of the household itself may respond to the OAP. Posel, Fairburn, and Lund (2004) report that adult co-residence is correlated with pension eligibility. Edmonds, Mammen, and Miller (2005) report that households with pension-eligible women tend to have more young children, more women in their early twenties, and fewer women in their thirties. Of course, if household composition only adjusts in terms of residents' standard measured characteristics (like age and gender), then the interpretation of other empirical findings at the household level is not affected so long as these characteristics are controlled.

In this paper, however, we present evidence that adjustments in household composition associated with the pension are *not* solely in terms of these characteristics, but are also related to measured and unmeasured individual characteristics which are themselves related to individuals' "bargaining power," their tastes, and therefore to a broad array of outcomes. Specifically, we observe associations between the OAP, the gender of the receipient, household demographic composition, *and heights and educational attainment of adults* in the household. Since adult heights and schooling were fixed by the time the OAP became available, we attribute these differences to a relationship between migration patterns, individual endowments, and the pension income. These results suggest that the oft-noted but heretofore little explored issue of household composition may have significant implications for the interpretation of the empirical literature on the OAP. To take one example, consider the finding reported by Bertrand, Mullainathan and Miller (2003) of a negative as-

sociation between prime-age adult labor supply and the pension income. In the absence of differential migration patterns, these results would imply that the pension income causes individuals who would otherwise work for a wage to allocate their time to leisure; however, once the possibility of differential migration is taken into account, possible alternative interpretations present themselves. One such interpretation is that individuals with a relatively lower market value of time (who would have been unlikely to be employed, regardless of the pension income) *move in* with the pensioner–possibly in order to provide services that the pensioner is able to demand. The implications, however, are more general than any single paper, and in fact are relevant beyond the context of the OAP.

The next section provides a brief overview of the South African OAP, highlighting the features of the program which recommend it as a vehicle for the analysis of family decision making. Based on these features, the following section discusses two broad empirical strategies for isolating effects of the OAP. Obviously, households which take up the pension are likely to differ from those which do not, and so a direct comparison of recipient to non-recipient households would not be useful. Following the literature, we exploit the age eligibility rule to identify the causal effect of the pension. However, we observe that despite its generosity, take-up of the pension by eligibles is not universal, and occasionally individuals who are still short of the age of eligibility report receipt of pension income. Therefore, in addition to the "intent to treat" type of analysis which has been common in the literature, in which age eligibility is treated as if perfectly predicting pension receipt, we also explicitly model takeup in an instrumental variables framework using the age/gender-eligibility rule as the instrument.

Next, we outline a conceptual framework for our analysis. We sketch a model of decision making at the level of the extended *family* (as opposed to the single household, which is a sub-family unit). The family allocates economic resources including not only wealth and market goods but also the time and services of individual members. In part, it does so by distributing human resources among multiple households. Our

exposition suggests some possible mechanisms of association of the OAP *not only* with household demographic composition *but also* with measured and unmeasured characteristics of each individual family member. Of course, these principles are not new–and in fact have been mentioned elsewhere in the OAP literature. However, our aim with the exposition is to reemphasize the empirical importance of issues which arise specifically in the context of sorting of individuals *on the basis of often unmeasured personal characteristics*.

Having framed the discussion, we proceed to the empirical analysis. First, using household level consumption data, a basic implication of the unitary model is tested: does pension income have an independent effect on household budget allocation, even after controlling total resources? We find that it does, but that this apparent deviation from "unitary" behavior does not conform easily with previous observations elsewhere in the literature. Specifically, we observe an association between pension income *when it accrues specifically to men* and a shift household allocation toward goods assignable to girls, and toward health care. This observation seems gently at odds with observations elsewhere in the literature which attribute improved outcomes in girls to pension income accruing specifically to women. As we discuss, the possibility that individuals are sorted on the basis of tastes or other personal characteristics admits of alternative–arguably more plausible–explanations.

Next, we explore directly the effects of the pension on living arrangements. In keeping with previous findings in the literature, we provide evidence of an association between the OAP and the demographic composition of the household. Then, we present evidence that OAP eligibility is associated not only with changes in the distribution of age and gender within the household, but also in the distribution of human capital–as indicated by the educational attainment and height of adult members. Since education and height are fixed in adults, these changes reflect the selectivity of those who coreside with OAP eligibles and OAP recipients. Generally, if these characteristics–and all the other characteristics associated with this selectivity–are not controlled in empirical models that seek to measure the causal effect of the

OAP, the estimates will be biased. We then emulate tests of the effects of the OAP which have been presented elsewhere in the literature, and attempt to extend them to explore the implications of individual sorting. Consistent with other findings in the literature, we find that both the presence of a pension eligible person in the household and pension income are positively associated with children's educational attainment. Furthermore, also consistent with results elsewhere in the literature, we find evidence that treatment effects differ based on the gender of the pension recipient and the gender of the child. We explore the possibility that some of these differences may be driven by differences in innate human capital endowments. The results are suggestive that the sorting of family members to households may play an empirically important role. We conclude with a discussion of the broader implications.

## Brief Overview of the South African OAP

The OAP is an extremely generous non-contributory public transfer of resources to older men and women in South Africa. The pension program was introduced in 1928 as an income supplement for elderly low income whites. In 1944, eligibility was extended to the "colored" and black population groups, with payment schedules dependent upon the recipient's race and location of residence. In 1965, differentiation based on residence was abolished but the racial gaps were maintained so that in 1975, for example, eligible whites received over R1200 (in 2004 prices), while coloreds received about R600 and blacks received less than R200. The racial gaps were reduced during the 1970s and 1980s and, with the fall of apartheid in the early 1990s, payments became race-blind. In 1993, the payment was set at R370 (which is R740 in 2004 prices) and the payment has remained approximately constant in real terms since then.

All women age 60 and older and all men age 65 and older who satisfy a very generous income and assets means test are eligible for the pension. In 2004, for example,

a single person was eligible for the grant if his or her monthly income was R1502 or less and if his or her assets were less than R266,400 (excluding owner-occupied housing). Only the income or assets of the pension recipient are taken into consideration; few blacks have income or assets above these levels. Accordingly, about 80% of age eligible black females and about 75% of males reported pension income in the year 2000. Nor is the amount of the benefit dependent on individual labor market participation or other endogenous factors. Individuals are eligible to collect the pension whether or not they work, and although the size of the benefit is supposed to vary according to the means of the recipient, among those who reported any pension income in the year 2000, over 90% reported receiving the maximum allowable payment of R740 (or about US$120). This benefit level is very high–about equal to the median income of black 20-50 year old income earners in South Africa, and about the 70th percentile in rural areas. In the year 2000, the OAP was the only income source in the median recipient black household in rural areas, and represented about two thirds of income in the median recipient black household in urban areas. Due to the age eligibility rule, pension benefits differ by gender; since a male who survives to age 65 can expect to live an additional 7 years and a female who survives to age 60 can expect to live an additional 10 years, the expected present discounted value of the pension is greater for female recipients than males. In short, the OAP is an extremely generous transfer to elderly blacks–and particularly to elderly black women–which is unrelated to these individuals' past decisions and which is primarily determined by their age at last birthday.

## OAP Takeup and the Magnitude of Causal Effects

Of course, *receipt* of the pension does not solely depend on eligibility, but is also associated with personal characteristics. For example, individuals must apply for the pension and demonstrate eligibility. Less patient individuals, therefore, may be less likely to pay the up-front cost of this effort in return for the discounted stream of income. Therefore, since pension *eligibility* for black South Africans–unlike pension

*receipt*–is determined almost entirely on the basis of two observable chracteristics (gender and age at last birthday), studies in the literature have examined the impact of eligibility on outcomes and thereby measured an "intent to treat" type treatment effect. However, despite its generosity takeup of the pension is not universal. If a fraction $\tau$ of eligibles take up the pension, and the impact on those who do take up is the same as the impact on the $(1-\tau)$ proportion who do not, then the estimated effect using the ITT approach can be interpreted as $\frac{1}{\tau}$ times the effect which the pension would have had if takeup were universal. However, if the decision to takeup is related to the effect of treatment, then it is likely to be important to model this decision explicitly.

Therefore, an alternative, more structural approach measures the impact of receiving pensions on the outcomes of interest, but models receipt as a choice in an instrumental variables framework. Age-eligibility is a valid and highly powerful instrument for pension receipt. It is excludable since it is determined almost entirely by two observable characteristics–age and gender. After controlling semiparametrically for these characteristics, therefore, it should have no independent relationship with any behaviors or outcomes other than through bargaining power and income. It is exogenous because the discrete jump from non-age-eligibility to age-eligibility is determined solely by an individual's date of birth, which is almost surely unrelated to any relevant personal characteristics.

Furthermore, the strength of the instrument is illustrated in figure 1, which shows the relationship between pension receipt, age, and gender in 2000 as reported in the South African Income and Expenditure Survey which is conducted every five years, and is used to compute the weights for the consumer price index (Statistics South Africa, 2000). Among blacks in particular, pension uptake jumps dramatically at the age of eligibility. This is also reflected in table 1, which shows pension takeup within 5-year age bands. Among black women, for example, takeup increases by 52 percentage points for those who are just age eligible relative to those who are almost age eligible.

However, another point highlighted by Table 1 is that eligibility is not a *perfect* predictor of pension receipt. For example, 15 percent of black men who are just short of age eligibility nonetheless report some pension income, and pension receipt continues to increase with age even beyond the age eligibilty threshold. Early takeup is particularly common among men, which is part of the reason that the presence of a male pension eligible individual in a household is a relatively weaker instrument for the presence of a male pension recipient.[1] Furthermore, in addition to early takeup, substantial *delays* in takeup are also observed. For example, only 60 percent of households with just-eligible individuals report receipt of pension income.[2]

The probability of takeup conditional on eligibility status is likely to be related to individual and household characteristics–for example, we observe in the IES that

[1]This early takeup may simply reflect misattribution of the pension income. For example, a man may report his wife's pension income as his own. However, if this were the case then the presence of other age eligible individuals in the household should be associated with the probability of early takeup. It is not. More likely, the pattern reflects less strict enforcement of the age eligibility criterion for men than women. Note in this context that both the intent-to-treat and instrumental variables estimates rely on the assumption that the pension has no causal effects on behavior before it becomes available, so that any behavior which is directly related to the pension will change discontinuously at the age of eligibility. But if it is in fact true that men take up the program more smoothly across the threshold of age eligibility, then inference based on such a regression discontinuity approach will understate the effect of the pension on men's behavior. More generally, *any* inference relies on a host of assumptions about the timing of onset of program effects, including the assumptions that older adults who are approaching the age of eligibility are not able to borrow against the promise of their future pensions, and that individuals' bargaining power within the household does not change because they are known to be approaching age eligibility. For example, if families with older adults run down assets in anticipation of the future pension or save less than they would without the pension, then behavior and outcomes would be observed to change smoothly around the age of eligibility, and the regression discontinuity approach will yield an attenuated estimate of the causal effect of the pension. This applies to both the intent-to-treat and instrumental variables estimates. In fact, the assumption of no significant smoothing across the age-eligibility threshold is fundamental to the empirical strategy employed throughout the literature, although it has little theoretical appeal.

[2]This may simply reflect reporting error, for example if substantial numbers of age eligible individuals receive pension income but fail to report it in the Income and Expenditure Survey. However, this would imply that these individuals are underreporting income; assuming that the other household respondents do not overreport income or underreport expenditure by exactly the same amount, these households should be expected to appear to have especially high levels of dissaving. However, savings rates in households that have a member who is eligible for the pension but does not report receipt of it are no different from savings rates in other households. Alternatively, it may be that the older adult is misreporting the OAP as some other source of income. Recently eligible respondents, for example, may confuse the names of other public transfer programs with the name of the OAP. If this is the case, non-OAP public transfers (and therefore total income) of households with individuals who just aged into eligibility but who do not report receiving the pension should be higher than in households with individuals about to age into eligibility. Neither implication is true. It is more likely that in fact take up is not universal.

African households with at least one eligible individual are more likely to report pension income if they are are larger, and if they are located in poorer provinces or rural areas. Therefore it is difficult to infer what the effects of the program would be if takeup were universal based solely on an intent to treat analysis. Accordingly, in the next section, when we analyze the relationship between pension income and the allocation of household resources, the discussion focuses on the results of instrumental variables estimates. However, since the majority of the literature on the OAP has reported ITT estimates, we also report on these results for comparability. Qualitatively, the results are similar.

## Conceptual Framework

In a family that can be characterized as having unitary preferences, the decision making unit is the family as a whole, so that any income transfer into the family should affect observed family allocations in exactly the same way, regardless of the individual to whom it is attributed–therefore, any behavioral effect of the OAP should operate solely through on income effect. In that case, the welfare impact of the OAP on every individual in the recipient families would be the same as that of any other income transfer program–and would be determined by the family's immutable common preferences. In non-unitary models of the family, however, family behavior is an aggregaton of individual level behavior, so that the impact of a transfer on allocation (or individual welfare) is less clear. Transferring income using the OAP would have a different effect from other income transfer schemes. In that case, characterizing the effect of the OAP on the welfare of individual family members becomes a question of even greater public policy interest. Therefore, most of the empirical literature on the OAP relates to the question of how well family behavior can be described by a "unitary" model (and relatedly, the particular welfare outcomes which derive from transferring income which is directly attributable to elderly men or women).

However, these tests are conducted at the level of the household, rather than the family. If families sort *themselves* into households (rather than being exogenously assigned), then empirical results must be treated with care. If the sorting of individuals to households is based solely on measured characteristics of the individuals, then any empirical test which controls properly for these characteristics is straightforward to interpret. However, if the sorting also depends on characteristics which are not typically measured–for example, tastes or innate endowments–then to the extent that these characteristics have an independent effect on the outcome of interest, estimates of the effect of the OAP will be biased. Intuitively, it is not difficult to imagine stylized examples of optimal sorting which might depend on these characteristics. A family may choose to group all the members who prefer listening to loud music into one household, and all those who prefer to read in another. If an income transfer (say, to some member of the reading household) increases the family's total demand for children's education, then it may be used to finance the move of a child who is not in school from the loud music household into the reading household so that he or she can begin his or her studies. An analysis at the *household* level will associate a *decline* in average children's education with the income transfer (unless, of course, tastes for loud music and reading are somehow included in the regression). Note, by contrast, that if the income transfer instead had gone to some member of the loud music household, then it would have been associated with an *increase* in average children's education.[3]

As we highlighted in the previous section, an important feature of the design of the South African OAP is that means testing does not take into account the incomes of anyone other than the potential recipient him or herself, whereas means-tested programs typically condition eligibility on resources of the household in which a potential recipient resides. Thus, in South Africa, unlike in the typical case, older adults have no incentive to change their living arrangements in order to qualify for the

[3]Finally, note that if the public transfer is targeted at the sorts of individuals who are more likely to live in one or the other household, then it may appear to have effects beyond the direct income effect, even if it does not. Controlling for total household income (but not total family income), one compares recipient (reading) households to other households with the same level of income (including both reading and loud music households).

pension income. However, there may be other reasons that pension receipt or pension eligibility could be related to household assignment. The assumption in much of the literature that "the pension program. . . provides no incentives to partition the household" (Duflo, 2003) may not be valid. Indeed, in the United States and Europe, pension income has been shown to enable adults to remain living independently at older ages (see, for example, Englehardt, Gruber and Perry, 2002). A less stylized example of sorting which might have empirical implications–and one which is more directly applicable to the OAP–would incorporate the "technology" of home production. Imagine that a family can transfer material goods at relatively low cost among its members, regardless of coresidence patterns (for example, through cash remittances), but that the most straightforward way to transfer individuals' time or services is through the adjustment of coresidence patterns themselves. If the pension income causes an increase in the pensioner's demand for services which are relatively labor intensive, then it will likely cause an adjustment in coresidence patterns. The labor for the increased home production will likely be provided by the individuals whose market wage is relatively lower. Therefore, the probability will increase that the relatively low human capital adult family members will coreside with the pensioner. Since data are only observed at the level of the household, any outcome which is related to the human capital of the adult family members will appear to be directly caused by the pension. Since these changes are occurring only in recipient households, then the households would appear to exhibit "non-unitary" behavior, *even* if the family's behavior could be described with a unitary model.

Depending on the nature of the "technologies" of home production, the pension may even affect the ways in which children are sorted on their personal characteristics. For example, it may be that some children's observed or unobserved endowments are such that production of their human capital is relatively more intensive in family members' labor, whereas the production of others' human capital is relatively more intensive in market inputs–for example, some children develop more valuable skills by doing household chores under the guidance of their grandmothers, whereas others develop more valuable skills by attending school or by working

in a family business. Then, if all children's human capital is normal in the family welfare function, the pension will increase the likelihood that these differently endowed children will be sorted to different households. Thus, the pension income would appear to "cause" observed differences in any characteristics which are related to these endowments. Finally, note that *a priori* it is impossible to make categorical statements about the relationships between the equilibrium assignment of individuals to households and individual characteristics unless one makes strong assumptions about the forms of the family welfare function, about individual preferences, and about the nature of the "technologies" of home production.

A model which incorporated these characteristics might follow the following sketch. Consider a family of $F$ individuals. Each of these individuals has distinct preferences represented by the utility function $\nu^f$. These preferences are conditioned on measured characteristics of all family members and unmeasured characteristics of the individual (including his or her idiosyncratic tastes), and are defined over the consumption of leisure and $R$ different home-produced goods and services. These goods/services, in turn, would be produced using $S$ different market inputs and also household members' time and human capital. The simplest model would have the market inputs be "tradeable" (transferable between households), and individuals' labor "non-tradeable" (say, only transferable via the adjustment of living arrangements). A more fully developed model would express transport costs for both as a function of the geographical location of every household–but the costs of transporting labor services would of course be higher.[4] The family would allocate individuals and transferable inputs across households and also leisure and consumption across individuals in order to optimize some function which aggregates individual preferences, taking into account time, resource, and "technological" constraints.

In the simplest case, the resource constraint would be common to the family–all income, regardless of where it accrues, could be used to provide market inputs, regardless of where they are employed. However, the time and "technological" con-

[4] Accordingly, such a model would have to take on the additional burden of endogenizing not only the assignment of individuals to households, but also the geographical location of each household.

straints would be specific to each household–production of final goods could only be done in the household where it is consumed, by combining market inputs (say, food) with the time and skills of resident household members (say, time and effort spent cooking). If the household and the family are one and the same, then this simply describes a traditional household model. However, if the family is allowed to choose how to partition individual members into households, then the equilibrium assignment will almost surely be related to a host of individual and family characteristics which are difficult or nearly impossible to measure.

A more formal exposition of this model is included in an appendix.

## The OAP and Allocation of the Household Budget

In this section we present a brief analysis of the allocation of the household budget which focuses on demand for goods associated with human capital investments, for "child" and "adult" goods, for "male" and "female" goods, and also for food. Using data for all black households included in the Income and Expenditure Survey, we estimated a series of Working-Leser Engel curves, with the share of the budget spent on each commodity group as the dependent variable in a multivariate regression model. The empirical model controls (the natural logarithm of) total household per capita expenditure with splines that have knots at the quartiles of $\ln(PCE)$. In addition, the regression includes the number of males and, separately, the number of females in the household in 5-year age groups (from 0-5 through 76-80). Magisterial district fixed effects are included to absorb variation in prices, labor market opportunities, and other community characteristics that might affect expenditure patterns. (A magisterial district is similar to a county in the United States.) Since the IV and ITT results are similar, in the interest of conciseness we present only the IV. Regression results are reported in Table 2. Estimates for three covariates are reported: whether there is a male who received the pension in the household, a female who received the pension, or whether there is both a male and a female who

received the pension.

The regressions control household resources (log of per capita expenditure) and the demographic composition of the household in a very flexible way. The impact of pension eligibility or pension receipt can be interpreted as being over and above the direct income effect–and so a unitary model of the family, together with an assumption that the selection of indivdiuals into households is independent of unmeasured factors which are themselves associated with demands, would predict that they would be indistinguishable from zero.[5] Alternatively, in a "non-unitary" model of the family, they would reflect the impact of increased decision-making power in the hands of the elderly.

The results in the table indicate that the pension income does indeed shift household budget allocations. Specifically, female pension income is associated with a 6.3 percentage point increase in the household's budget allocation to food, a 0.9 percentage point decrease in the allocation to male clothing, and a 0.8 percentage point decrease in the allocation to tobacco. Furthermore, about a third of the shift toward food expenditures in these households was accounted for by vegetables (1.3 percentage points) and dairy (0.9 percentage points) (in the interest of conciseness, these results are not reported in the table). These may be the ingredients of a diet which is healthier for children.[6] These results suggest a picture of grandmothers allocating more

[5]There are several other reasons why estimated effects of pension income might be different from the effects of other resources. First, changes in budget allocations may reflect changes in tastes over the life course that are unrelated to pension income. However, such taste changes are unlikely to be discrete and occur at age 60 for females but age 65 for males. Second, there may be non-linearities in the effects of resources which are not fully captured in the empirical Engel curves. A specific type of non-linearity relates to variance of income. It may be that pension income is treated differently from other income because it is less variable over time than other sources of income. In these cases, the sources of pension income (whether from pensions paid to males, females, or both) should have similar effect on demand–though these effects may be scaled since the expected time horizon of the payments to males are smaller than to females. This suggests comparing the three estimated income effects.

[6]However, note that pension income which goes solely to female recipients is also associated with a significant 0.4 percentage point increase in the household budget share for fats, while in households with only male recipients and with both male and female recipients, the pension income is associated with a shift *away* from fats (of -0.9 percentage points and -0.6 percentage points, respectively–although only the former is statistically significant at the 5% level). Given the high levels of obesity in South Africa, particularly among older women (Pouande and others, 2002), it is not obvious that this difference reflects an association between higher quality diet and having a female pensioner in

resources toward goods and services which are good for younger children, and so are consistent with others reported in the literature in support of the conclusion that pensions paid to women tend to result in better outcomes for grandchildren (Case and Deaton) and the longer-term nutritional status of granddaughters (Duflo).

However, the results also seem to suggest that it is only *male* pensioners–and not female pensioners–who shift household resources toward consumption of girls' clothing. This may suggest grandfathers are relatively more altruistic toward their granddaughters. This interpretation, of course, seems to be gently at odds with the stylized fact that improving the bargaining power of grand*mothers* improves the care received by granddaughters. Similarly, male pension income is associated with a 1.4 percentage point increase in expenditure on health care, and male and female pension income together are associated with an even larger 2.3 percentage point increase, which may suggest that elderly men are relatively more inclined to increase the household's consumption of health care. This too would not align neatly with the observation that the health of at least some household members is improved when the pension income accrues to elderly women.

If the composition of the household itself may be determined by the pension, other (arguably more plausible) interpretations may . The regressions control flexibly for the effect of the pension on the *age* and *gender* of household members, but the pension may affect other characteristics. For example, the increased budget share for girls' clothing which is associated with male pensioners may not be a direct result of grandfathers' relative altruism toward their granddaughters at all. It may be that the granddaughters who coreside with male pensioners differ from those who do not in ways which are not measured in these data. For example, if the pension accruing to a man in a family increased the probability that he would coreside specifically with his oldest granddaughter, then the lower expenditure on girls' clothing in non-recipient or female recipient households would simply be accounted for by the fact that girls in those households could inherit the clothing of their older sisters. With regard to the expenditures on health care, it is always perilous to use *household* bud-

the household.

get data to draw conclusions about *individuals'* consumption, but the interpretation of this finding is *especially* dependent on the distribution of individual characteristics within the household. On the one hand, this may reflect spending for the health care of older adults like the pensioners themselves.[7] In that case, the fact that female pensioners do not skew the allocation of the budget toward health care might be explained by the presence of family members with a comparative advantage in caring for them at home. As we show below, the adult women who coreside with female pensioners are those likely to earn lower wages in the labor market. Keeping them home to care for the pensioner would allow the household to substitute in-home labor for purchased health services. Alternatively, if male pension income attracts or retains individuals who are more likely to be (or become) ill, this could also account for the observed differences.[8]

## Pensions and Household Living Arrangements

The model of family behavior which we have sketched highlights some of the mechanisms by which living arrangements may be endogenously determined–and more importantly, by which coresidence patterns might vary based on observed and unobserved characteristics of family members. In table 3 we present empirical evidence of an association between living arrangements and pension eligibility based on a comparison of the demographic composition in pension eligible and non-eligible households. Data are drawn from the Income and Expenditure Survey. We assemble a vector of eight dependent variables, consisting of the number of individuals resident in respondent households aged 0-40 in each 10-year age/gender band. Using seemingly unrelated regression, we simultaneously regress each element in the dependent vector against a complete set of demographic controls (that is, the number

[7]Note that this might represent an expansion of the family's Pareto frontier, if the spending keeps pensioners or future pensioners in better health so that they live longer and thus the pension income is paid for a longer period of time

[8]Below, we show that relatively less educated males are those most likely to be coresiding with male pensioners. If these men were particularly likely to be showing symptoms of HIV infection at the time of the 2000 IES, this might explain the diversion of the household budget toward health care.

of individuals in all the other age/gender bands), and a pair of dummy variables which test for a separate effect which is common to all pension eligibles and female pension eligibles (thus, the reported results are intent-to-treat). The SUR approach allows us to exploit cross-equation correlations to improve the efficiency of our estimates, and also to compare coefficient estimates across equations. The results in Table 3 indicate that households with sole female pension eligibles are more likely to contain older boys and younger girls, and less likely to contain adult women aged 20-50. A test for equality in the relationship between pension eligibility and the presence of women in each of the ten year age bands (20-29, 30-39, 40-49) fails to reject the hypothesis that the effects are equal. These findings are consistent for example with those reported by Posel, Fairburn, and Lund, who associate the OAP with the departure of prime-age women from the household.

It can hardly be surprising that living arrangements respond to the pension, given its magnitude. However, if these responses involve changes in household composition that are uncorrelated with other characteristics that might affect consumption behavior or human capital outcomes, then the effects of the pension on these outcomes is straightforward to assess based on household-level data. Table 4 reports the relationship between adult human capital indicators and the household's receipt of pension income. The data are from the South African Demographic and Health Survey (DHS). This survey was conducted in 1998 and is representative of all households throughout the country (Measure DHS, 1998); it collected detailed demographic and reproductive histories from over 11,000 women age 15 through 49 in conjunction with socio-demographic characteristics of all household members and measures of health, including anthropometry, for all household members age 15 and older.

For the regressions in table 3, human capital is measured by adult education (in years of schooling) in panel A and adult height (in cms) in panel B for adults age 20 through 55. Since these outcomes were fixed when the OAP became available, we interpret any relationship between these indicators and household pension receipt

as a reflection of selectivity in coresidence patterns. In addition to pension receipt in the household, the regressions control age of the respondent (with dummies for each five-year age band), household composition (numbers of males and females in each five year age band) and household wealth (indicators of certain assets). The analyses are stratified by gender. Instrumental variable estimates are reported with pension eligibility of one or more household members as the instruments.[9]

The first panel reports the relationship between adult human capital and whether anyone in the household received a pension–regardless of the gender of the recipient. The results indicate that shorter prime-age females are likely to coreside with pensioners. Specifically, a household's receipt of pension income is associated with about a 4 cm height disadvantage for resident adult women, relative to women in households not receiving the pension. (In addition poorer-educated adults and taller males may be more likely to reside in pension-recipient households, although these relationships are not statistically significant). Perhaps the pensioner's improved "bargaining power" shifts the family's consumption toward home-produced goods or services which these individuals have a comparative advantage in producing (and which are produced and consumed, of course, in the pensioner's household). For example, if the elderly prefer the sort of care which a family member can provide most efficiently[10], and if shorter women are also those who have the lowest wage prospects, an elderly man or woman may be able to effectively "hire in" his or her shortest adult daughter to provide it. Alternatively, perhaps the pension income causes a shift in family time allocation which involves relatively taller females outmigrating from pensioner households. A fully developed model of family decision making would likely admit of several of mechanisms by which this outcome might emerge, and we do not presume to have enough evidence to distinguish between them. Rather, throughout this section, by highlighting a few alternative mechanisms which could generate the observed outcomes, we only mean to emphasize the complexity of interpreting household level analysis

[9] In the interest of conciseness, we do not report ITT results here; they are substantially the same.
[10] For example, cooking, cleaning, in-home nursing, or even simply company.

when family members are likely sorted on unobservable characteristics.

The second panel explores whether the gender of the recipient affects the observed selectivity.[11] The top part of the panel considers all adults age 20-55 together. The pattern of shorter women coresiding with pensioners appears to be irrespective of the gender of the pensioner (although the effect is more precisely estimated for women in households with female pensioners). However, the selectivity of males does appear to depend on the gender of the pension recipient. Taller and less well-educated adult men are significantly more likely to live in a household with solely *male* pensioners. Pension income accruing solely to males is associated with a 1.7 year education "disadvantage" and a 4.6 cm height "advantage" for adult males. By contrast, adult men who live with female pensioners are not discernibly different in terms of either height or education from those who live in households without pension income (although the difference in the effect on height is not statistically significant). One of the many possible mechanisms which could account for this outcome would involve males using the pension income to start or expand family enterprises. If male adults with better innate health endowments (reflected in their heights) but worse market wage prospects (reflected in their education) have a comparative advantage in this type of production, and if their skills are complementary with the male pensioner's managerial skills, then it may be most efficient for the family to colocate these individuals. Alternatively, it may be that the tastes of male pensioners for household public goods are similar to the tastes of poorer educated but healthier male adults, and that the pension income shifts the family's consumption toward these public goods.[12]

In the bottom part of the panel, we analyze effects on younger adults (20-30) sep-

[11] We grouped the 680 individuals in households with both male and female pensioners together with the 2985 individuals in households with only female pensioners, because in analyses not reported here, the treatment effects in these two groups were found to be similar. We also tested whether the *amount* of pension income matters, separating the relationship in households with one pensioner from those with more than one pensioner. Cell sizes were too small to detect significant differences and so we did not pursue this line of analysis.

[12] A stylized example might be the following: imagine the pensioner uses his income to support his household's regular reception of broadcasts of sporting matches, and healthier but less well-educated males are most likely to share his interest in sports.

arately from older (31-55). The negative selection on education of male adults by male pensioners is observed in both age groups and is of comparable magnitude (-1.5 years for male adults aged 20-30, -1.9 years for male adults aged 31-55). Furthermore, the *difference* in the effects of male and female pension income is significant for both age groups, although possibly for subtly different reasons. Specifically, for *younger* adult males, the evidence points to a *positive* association of education with female pension income (and a negative or zero assocation with male pension income). By contrast, for *older* adult males, the evidence indicates a *negative* assocation between education and male pension income (and a positive or zero association with female pension income). Referring to the example above, it may be that older adult males who are more poorly educated are the ones with the greatest comparative advantage in home enterprise production, which would account for their more discernible positive selection on education, or that the tastes of male pensioners are more likely to resemble those of older, more poorly educated adult males than younger, more poorly educated ones. To highlight another layer of complexity, note that there is evidence of a negative association between male pension receipt and the education of younger adult *females* (-1.4 years, significant at the 10% level). Therefore, the observed effect of male pension income on education levels of either males or females may simply be a reflection of assortative mating and differences in the probability of coresidence among adult children, parents, and their spouses.[13] Similarly, female pensioners may retain or attract daughters who are of higher "quality" in some unmeasured way, so that it is their relatively better educated sons-in-law who account for the positive association between female pension income and adult male education.[14]

The bottom part of the panel also sheds light on the negative assocation between

[13]Living arrangements in South Africa are extremely complex. Many couples live apart, with men working in a town and sending remittances to their wives and children who live in a rural area. Evidence from longitudinal data suggests that between 1993 and 1998, a significant fraction of women moved from rural to urban areas in order to be with their husbands and to work. Their children stayed in the rural areas with the grandmother. (Maluccio and Thomas, 2005).

[14]Alternatively, female pensioners may simply prefer the company of their better educated sons, and so use the pension income to retain them, or to finance the outmigration of those men who are more poorly educated.

pension income and the heights of females. As discussed above, the *average* overall effect of pension income on female adult heights is the same irrespective of the gender of the recipient. However, female pension income is associated with an enormous (-8 cm) height gap among younger adult females (aged 20-30), while male pension income is associated with a relatively smaller, statistically insignificant height gap. By contrast, for older adult females the large negative association (of -5.5 cm) is with male pension income, whereas it is female pension income for this group which is associated with a smaller, stastically insigificant gap.

In short, these results present a picture of household composition responses to the pension which are complex and which differ not only based on the gender of the recipient, but also based on the gender, age, and levels of fixed human capital of individuals, and on interactions among all of these factors. In the following section, we revisit a previous pattern observed in the literature on the OAP, highlighting some of the implications that these responses might have on the results.

## Pensions and Child Educational Attainment

This section examines the relationship between pension receipt and the educational attainment of all 6 to 19 year olds in the household. Attainment is measured as the level and grade in which the respondent is currently enrolled (or was last enrolled) translated into the number of years of schooling necessary to attain that grade. Data are drawn from the DHS.

Multivariate regression results are reported in table 4. The regressions are estimated separately for boys and girls and control age of the child (with an indicator for each year of age), household composition (number of males and number of females in 5-year age groups), and indicators for ownership of a set of assets. The first panel of the table reports estimates of the relationship between educational attainment and living in a household with an age-eligible person (that is, the intent-to-treat effects). They are indistinguishable from zero. The second panel reports instrumental

variable estimates of the effect of living in a household with a pensioner (with age eligibility as the instruments). The results parallel the intent-to-treat results. In the rest of this section, we discuss only the IV results.

It is possible that no effects are detectable since the relationship between educational attainment and pensions differs depending on the age of the child. Young children may delay entry into school because of liquidity constraints or because the children are providing services at home. In that case, pension income may relieve those constraints and the children will attend school. Older children, by contrast, may leave school for a different set of reasons, and so the effect of the pension on their outcomes may be distinct. Accordingly, we stratified the samples into young children (age 6 through 12) and older children (age 13 through 19). Pension income is associated with lower school attainment among older children and among young boys (although none of these coefficients is significant) but has a significant positive relationship with the education of young girls (at a 10% size of test). Moreover, the difference between the estimate for boys and girls is significant (the estimated difference of 0.675 additional years' worth of treatment effects on girls relative to boys has a standard error of 0.271) indicating that, in terms of educational attainment, a pensioner in the household benefits young girls more than young boys.

This is consistent with the evidence reported by Duflo who found that in households with a pensioner, girls have better height-for-age and weight-for-height indicators, whereas boys do not. She also reported that this effect was greater if the pensioner was a woman. Accordingly, in Panel B we allow for separate effects of pension receipt depending on the gender of the recipients.[15] In addition to pooled estimates for all children, we estimate effects on younger and older children separately. Among younger children, having a female pensioner in the household is associated with better educated girls and poorer educated boys (the estimated difference of 0.64 years has a standard error of 0.290).

[15]To improve the power of our tests, we grouped the 214 boys and 238 girls living in households with both male and female recipients together with the 254 boys and 220 girls living in households with only female pensioners, since (in results not reported here) the magnitudes of treatment effects were observed to be similar in these two groups.

However, if there is male pension income in the household, girls' educational attainment is even higher (1.5 years greater, albeit at a 10% significance level). No benefit accrues to boys. These results contrast sharply with those for older children, for whom male pension income is associated with higher levels of education among boys but lower education among girls. The impact of the recipient's gender on the effects of pension income on older boys as against older girls is large (3.61 years, with a standard error of 1.263). Consistent with these findings is the fact that even *within* households (that is, including a household fixed effect), the effect of male pension income is associated with a 0.33 year stronger treatment effect on girls than on boys (with a standard error of 0.15), whereas female pension income is associated with a 0.27 year more powerful treatment effect on girls than on boys (with a standard error of 0.10).

Taken together, the results in Table 4 present a picture which is generally consistent with the previous literature of the effect of pension income on children and adolescents (in that the effects of pension income depend on both the gender of the recipient and the gender of the child), but which is slightly more nuanced. Among older children, male pension income is associated with positive educational outcomes among boys and (significantly different) negative outcomes among girls, whereas female pension income has roughly the same negative or zero effect on both girls and boys. By contrast, among younger children, female pension income has a positive or zero effect on girls and a (significantly different) negative or zero effect on boys.

The complexity of the picture is not suprising, since as the previous section illustrates, the pension income is likely affecting *not only* investments in the children's human capital, but also the distribution of unmeasured endowments among the household members. For example, given the findings in the previous section and the fact that innate human capital endowments are correlated across generations, if shorter women who move in with a pensioner are likely to bring their children (or, observationally equivalent in these data, if taller women who depart from the

household are likely to take their children with them), an estimate of the causal relationship between pension income and measures of child human capital will be biased downwards (and might even appear to be negative). One way to examine whether these selection dynamics may also be driving the observed relationships between the pension and children's education is to compare the apparent treatment effects on children who are co-resident with their mothers against those who are not. Since women co-resident with pensioners are generally shorter and less educated, it is likely that the women who departed from the household as a result of the pension (or, equivalently, those in the extended family who did not join the household) are relatively more educated and taller. Therefore, in the absence of any direct information on these children's mothers, we estimate the difference in apparent treatment effects of the pension on children co-resident with their mothers relative to its apparent effects on those not co-resident with their mothers. Thus, this regression represents a difference-in-difference estimate. The results are presented in Table 5. The data are drawn from the DHS, and as in previous regressions, age eligibility instruments for pension receipt, the age of the child (in the form of dummies for each year of age), and a vector of household compositon indicators are included.

Focusing on column 1, we note that in general boys aged 6-15 who are coresident with their mothers are further ahead in school than those who are not. Furthermore, as shown in the previous results on child attainment, those living with a pensioner are no more or less educated than those who are not. However, the apparent treatment effect among children co-resident with their mothers is significantly *negative* relative to those who are not–specifically, boys co-resident with their mothers appear to benefit from the pension by a half year less than boys who are not co-resident with their mothers. This pattern is also observed for girls (column 4), although the magnitudes are smaller and the result is only significant at the 10% level. This pattern may reflect a complicated decision rule over how to allocate the benefits of pension income across grandchildren, but more likely, it reflects the fact that those children who are co-resident with their mothers are, like their mothers, relatively disadvantaged in terms of their innate human capital endowments. These results

suggest that the problems of selectivity in coresidence patterns associated with the pension complicate the inference of child outcomes, because those individuals who coreside with pensioners may differ in unmeasured ways from those who do not.

In order to explore this further, we examine data from the KwaZulu Natal Income Dynamics Survey (KIDS), a small longitudinal panel which reports income and assets as well as health status–including anthropometrics for certain household members–for about 1200 households in KwaZulu Natal province at two points in time–1993 and 1998.[16] Given the confounding effects of household sorting, one might prefer to assess the effects of the pension on child well-being using a difference-in-difference sort of approach, directly comparing *changes* in attainment for those who were exposed to pension income against those who were not. However, note that revisiting the same household at two points in time would not allow for such an analysis, since in order for the change in schooling to be measured in that case, a child would have had to reside in the household both times. If pension households tend to retain children who are less (or more) likely to continue in school for any reason, then the pension will appear to "cause" reduced (or increased) improvements.[17] If some children's endowments and prior schooling attainment are such that they would learn more valuable skills in home production than in further schooling, then the family may well choose to colocate them in the households with the pensioners–which, after all, are more likely to be specialized in home production. Children who would learn more valuable skills by staying in school would be more likely to be relocated elsewhere.

To explore this possibility, in table 6 we examine children's educational status in 1993 as a function of present *and future* availability of the pension income. The first

[16]The baseline sample consisted of 1400 households, of whom some 16% were lost to follow up for the 1998 interviews.

[17]Such an analysis using the KIDS data would relate negative or zero improvement in schooling attainment to the pension. Among boys aged 2-16 in 1993 who were observed again in 1998, those who were coresident with male pensioners had gained an average of 0.4 (standard error, 0.14) fewer years of schooling over those five years than those of the same age who were not coresident with pensioners, while those who were coresident with female pensioners gained 0.1 fewer years of schooling (standard error, 0.09). For girls, the observed relationships are also negative, though closer to zero and statistically insignificant.

column of the top panel indicates that children aged 6-19 who *would* five years later turn out to be coresident with a pension-eligible individual *had already* acquired a quarter of a year more schooling those of the same age and gender who would not turn out to be coresident with a pensioner. These results may in part reflect a failure of the assumption that households cannot change their current allocation of resources in anticipation of future income. However, the observed difference cannot be entirely explained as the result of a pension-induced shift in the allocation of household resources, because among those children who would turn out to be coresident with a pensioner in 1998, the ones who were *already* coresident with a pensioner in 1993 had nearly a third of a year *less* schooling. Therefore, the positive relationship between schooling and pension eligibility seems to be *negated* by extended exposure to the pension. These results suggest a complex set of relationships among living arrangements, children's education, and the pension income. For example, these results would obtain if the pension income increased the probability that a household would retain more educated children for a year or two before sending them off to better prospects elsewhere. They also would obtain if, in the absence of the pension, relatively less educated children are more likely to depart from a household, whereas the pension income reverses this relationship, accelerating the outmigration of relatively more educated children. This might be expected if the pension households are those which specialize in home production. There are many other dynamics which are empirically equivalent in these data.

# Conclusions

The OAP is a very generous cash transfer which provides each older man and woman in South Africa with a guaranteed income that places him or her around the median of the income distribution in the country. Eligibility for the OAP is determined by age and gender along with a means test which is binding for few black South Africans. In a series of very creative studies that explore decision-making within households, the OAP (or eligibility for it) has been treated as a source of

exogenous variation in income. In these studies, noting the interpretation of the results is predicated on the assumption that living arrangements do not respond to the OAP in ways that are not controlled, the authors highlight that this assumption is both necessary and warrants further investigation. This paper takes up that challenge.

Put simply, if family members are sorted into households on the basis of characteristics that are not controlled in the models, then comparisons of outcomes between pension eligible and non-eligible (or recipient and non-recipient) households capture not only the direct effects of the pension, but also the effects of these changes on living arrangements. The evidence marshaled in this paper demonstrates that living arrangements and, therefore, household composition does respond to eligibility for the OAP. To the extent that these responses are controlled in empirical models, the interpretation of that evidence remains unaffected. However, there is also evidence that adults with lower levels of human capital, as indicated by height and education, are more likely to co-reside with pension-eligible adults. Human capital of individuals is not typically controlled in empirical models of the impact of the OAP and so attributing estimated effects of the OAP to the impact of exogenous variation in income is potentially misleading.

An interpretation of the evidence presented above is that as older adults become eligible for the OAP, they demand more services of family members. These are provided by adult family members whose value of time is relatively low and who co-reside with the older adult because it is more efficient to provide the services the older adults wants if one co-resides with him or her. In the absence of data on the entire set of potential service-providers (the entire kin network), it is not possible to distinguish this interpretation from one in which the OAP is used to support prime age adults who choose to remain unemployed as suggested by Bertrand et al.

This research highlights the potential importance of moving beyond theory and data that is bound by the confines of a spatially-determined definition of a household to a broader conceptualization of family, kin and social networks that play an important

role in decisions about family life. Systematically including non co-resident kin in these models seems like a straightforward first step.

# References


Ardington, E. and F. Lund. 1995. "Pensions and Development: Social security as complementary to programmes of reconstruction and development." Development Southern Africa 12.4.

Bertrand, M., S. Mullainathan and D. Miller. 2003. "Public policy and extended families: Evidence from pensions in South Africa." World Bank Economic Review, 17.1:27-50.

Case, A. and A. Deaton. 1998. "Large cash transfers to the elderly in South Africa." Economic Journal, 108.450:1330-61.

Duflo, E. 2000. "Child health and household resources: Evidence from the South African Old Age Pension Program." American Economic Review: Papers and Proceedings, 90.2:393-8.

Duflo, E. 2003. "Grandmothers and Graddaughters: Old-Age Pensions and Intra-household Allocation in South Africa", World Bank Economic Review, 17:1-25.

Edmonds, E., K. Mammen and D. Miller. 2005. "Re-arranging the family?" Journal of Human Resources, 40.1:186-207.

Heckman, J., R. Lalonde and J. Smith. 1999. "The economics and econometrics of active labor market programs". Mimeo, University of Chicago.

Jensen, R. 2004. "Do private transfers displace the benefits of public transfers? Evidence form South Africa?" Journal of Public Economics, 88.1:89-112.

Legido-Quigley, H. 2003. "The South African Old Age Pension: Exploring the role on poverty alleviation in households affected by HIV/AIDS." Mimeo, University of Manchester.

Moffitt, R. 1991. "Program evaluation with non-experimental data." Evaluation Review, 15.3:291-314.

Posel, D., J. Fairburn and F. Lund. 2004. "Labour migration and households: A reconsideration of the effects of the social pension on labor supply in South Africa", mimeo, University of Kwa-Zulu Natal.

# Appendix

As indicated in the text, consider a family of $F$ individuals. Each of these individuals has distinct preferences represented by the utility function $\nu^f$. These preferences are conditioned on measured characteristics of all family members and unmeasured characteristics of the individual (including his or her idiosyncratic tastes), and are defined over the consumption of leisure and $R$ different home-produced goods and services. These goods/services, in turn, would be produced using $S$ different market inputs and also household members' time and human capital. For now, say that market inputs are "tradeable" between households–that is, they can be reallocated among individuals are very low cost–and individuals' time and skills are not–that is, they are for the most part transferred by adjusting coresidence patterns. Other notation is as follows:

- Similar in spirit to Becker's representation of the marriage market, the allocation of individuals to households can be denoted with the $F \times F$ matrix $\mathbf{A}$, where $a_{ij} = 1$ if person $i$ lives in a household headed by person $j$, and 0 otherwise.[18] Its $f$th column is denoted $\mathbf{a}^f$.

- The allocation of leisure time is denoted by the $F$-dimensional vector $\mathbf{t}^L$, where $t_f^L$ is the amount of leisure time consumed by individual $f$.

- The allocation of time spent in home production is denoted by the $F$-dimensional

[18] Assume that each individual must reside in one and only one household.

vectors $\mathbf{t}^r$, where $t_f^r$ is the amount of time devoted by individual $f$ to production of good or service $r$.

- The definitions of "leisure" and "household production" are such that any time not devoted to either of these two is necessarily devoted to wage labor.

- The allocation of goods and services within the family is denoted by the $F$-dimensional vector $\mathbf{z}^r$, where $z_f^r$ is the amount of good or service $r$ consumed by individual $f$.

- Human capital is reflected in relative skills in home and market production. These skills can be represented in terms of "relative effectiveness" of time, which are denoted by the $F$-dimensional row vectors $\psi^W$ and$\psi^r$. For example, $\frac{\psi_j^W}{\psi_i^W}$ denotes the number of hours devoted by family member $j$ to wage labor which is as productive as a single hour devoted by household member $i$. Therefore, normalizing $\psi_1^W$ to one, $\psi_f^W$ is the "effectiveness" of family member $f$ in wage labor. Similarly, $\psi_f^r$ represents the "relative effectiveness" of individual $f$ in household production of good or service $r$. For this exposition we are treating the $\psi$ parameters as if exogenously given, although in a fully developed model it would be under the control of the family. For example, the family's choice of geographic location for each household will determine the values of $\psi^W$, through differences in local prices. Also, in a dynamic model with learning-by-doing, the allocation of time in one period can affect the values of the $\psi$ variables in subsequent periods. Also, investments in children in the present could determine future values for the $\psi$ variables.

- The allocation of market inputs is denoted by the $S \times R$ matrix $\mathbf{K}^f$, where $k_{sr}^f$ is the amount of market input $s$ which is used to produce good $r$ in a household headed by individual $f$. If individual $f$ does not head a household, then $k_{sr}^f = 0$.

- The measured characteristics of the family which condition individual preferences are denoted by the matrix $\mathbf{M}$, and the unmeasured characteristics of

individual $f$ which condition his or her preferences are denoted by the vector $\varepsilon^f$.

The optimization problem therefore is:

$$W^\star = \max_{\substack{\mathbf{A}, \mathbf{z}^1, \ldots \mathbf{z}^R \\ \mathbf{t}^1, \ldots, \mathbf{t}^R, \mathbf{t}^L, \mathbf{K}^1, \ldots, \mathbf{K}^R}} \left\{ W \left[ \nu^1 \left( \mathbf{z}^1, \ldots \mathbf{z}^R, \mathbf{t}^L;\ \mathbf{M}, \varepsilon^1 \right), \ldots, \left( \mathbf{z}^1, \ldots \mathbf{z}^R, \mathbf{t}^L;\ \mathbf{M}, \varepsilon^F \right) \right] \right\}$$

subject to:

- First, the single family-wide resource constraint for market/transferable inputs, combined with the $F$ individual time constraints. (Denote the multiplier on this constraint $\alpha$):

$$y + \omega \bar{T} \sum_{f=1}^{F} \psi_f^W = \sum_{h=1}^{H} \sum_{r=1}^{R} \sum_{s=1}^{S} p_r k_{sr}^h + \omega \sum_{f=1}^{F} \psi_f^W \left( t_f^L + \sum_{r=1}^{R} t_f^r \right)$$

  where $y$ is the total family non-labor income, $\omega$ is the market wage for the individual of the family against whom the "relative effectiveness" of the others' time is normalized, $p_r$ is the price of market input $r$, and $\bar{T}$ is the fixed time endowment of 24 hours per day.

- Second, the $RF$ constraints imposed by the "technology" for the household production of final goods and services. (Denote the multipliers on these constraints $\beta_h^r$):

$$\left(\mathbf{z}^r\right)' \mathbf{a}^h = Z^r \left( k_{1r}^h, \ldots, k_{Sr}^h,\ \sum_{f=1}^{F} a_{fh} \psi_f^r t_f^r;\ \mathbf{M}, \xi \right) \quad \forall h \in \{1, \ldots, F\}, \forall r \in \{1, \ldots, R\}$$

  This technology is represented by the production function $Z^r(\cdot)$, which takes as arguments the amount of each market input and the total "effective time" commited within the household. The productivity of these factors is conditioned on measured and unmeasured ($\xi$) characteristics of the family, which

include the aptitude of each individual.

- Third, the $F$ constraints imposed by the fact that no individual can work a negative number of hours for a wage. (Denote the mulitipliers $\gamma_f$):

$$\bar{T} \geq t_f^L + \sum_{r=1}^{R} t_f^r \quad \forall f \in \{1, \ldots, F\}$$

Note that if every member of the family is exogenously assigned to the same household, headed by individual $h$, then only the total family resource constraint:

$$y + \omega \bar{T} \sum_{f=1}^{F} \psi_f^W = \sum_{r=1}^{R} \sum_{s=1}^{S} p_r k_{sr}^h + \omega \sum_{f=1}^{F} \psi_f^W \left( t_f^L + \sum_{r=1}^{R} t_f^r \right)$$

and the "technology" constraints become:

$$\sum_{f=1}^{F} z_f^r = Z^r \left( k_{1r}^h, \ldots, k_{Sr}^h, \sum_{f=1}^{F} \psi_f^r t_f^r; \mathbf{M}, \xi \right)$$

Solving the optimization problem in this special case entails distributing the total amount of each good in order to equalize the marginal utility of each family member's consumption of each good, weighted by the impact of each individual utility on overall family welfare (assuming interior solutions).

$$\sum_{f=1}^{F} \frac{\partial W}{\partial \nu^f} \cdot \frac{\partial \nu^f}{\partial z_i^r} = \sum_{f=1}^{F} \frac{\partial W}{\partial \nu^f} \cdot \frac{\partial \nu^f}{\partial z_j^r} \quad \forall i, j \in \{1, \ldots, F\}, \forall r \in \{1, \ldots, R\}$$

In turn, the weight given to member $f$'s preferences in the welfare function will depend on his or her power in family decision-making. Call the parameter reflecting this power $\lambda$. It depends on observed characteristics of all family members, wages and prices, non-labor income, and other factors that are associated with authority and influence of the individual member. Some of these might be measured characteristics, $\mu$, while others are, in terms of this model, unmeasured and might include altruism, patience, willingness to strike a hard bargain, tastes for risk, and so on.

These are denoted by $\zeta$.

$$\lambda = \lambda\left(\mathbf{p}, \omega, y, \psi^1, \ldots, \psi^R, \psi^W, \mathbf{M}, \mu, \zeta\right)$$

Total household demand for each market input and total household labor supply will be those which produce a bundle of household goods satisfying the condition above and the household's resource constraint as efficiently as possible. These demands, therefore, will depend on all prices, wages, non-labor income, and observed and unobserved characteristics, as well all the other factors which determine the distribution of bargaining power within the family. Combining demand for all the market inputs into the vector $\eta$, the demand function will be:

$$\eta = \eta\left(\mathbf{p}, \omega, y, \psi^X, \psi^Z, \psi^W, \lambda\left(\mathbf{p}, \omega, y, \psi^1, \ldots, \psi^R, \psi^W, \mathbf{m}, \mu, \zeta\right); \mathbf{M}, \varepsilon, \xi\right)$$

The reduced form will be $\eta' = \eta'\left(\mathbf{p}, \omega, y, \psi^1, \ldots, \psi^R, \psi^W, \mathbf{M}, \mu, \varepsilon, \xi, \zeta\right)$. Three points are highlighted by this exposition. First, if the family is unitary, then the effect of non-labor income on demand for market goods will be the same independent of the source of income and to whom the income is attributed. That is, demand is given by $\hat{\eta} = \hat{\eta}\left(\mathbf{p}, \omega, y, \mathbf{M}, \varepsilon, \xi, \zeta\right)$. Comparison of the function $\hat{\eta}$ with $\eta'$ yields two familiar tests for whether the family is unitary. First, pension income should have the same impact on demand as any other income. This will arise if the family/household welfare function $W[\cdot]$ aggregates individual utilities in some fixed way–for example, if the pensioner has the same preferences as all other family members or if all income is passed over to some other family member who then allocates resources. Second, under the assumptions of the unitary model, indicators of power, $\mu$, will not influence demand. In a non-unitary family, receipt of pension income will affect $\lambda$, and therefore demand. Moreover, eligibility alone may improve an individual's bargaining position within the family. Thus, eligibility may be one of the observable characteristics incorporated here in the vector $\mu$, which in turn will affect demand through $\lambda$.

However, interpreting evidence of "non-unitary" behavior is more difficult in a model where the **A** matrix is chosen. Since the $a_{fh}$ 's are omitted from any regression, any correlation among them, the variables $\psi^r, \psi^W, \mu$, and household demand for market inputs will produce biased estimates of the effects of these measures of "power." The exposition above, however, suggests that the equilibrium assignment will almost certainly depend on prices, wages, non-labor income (including the pension), human capital, and the other measured and unmeasured characteristics of household members, and also will depend on the forms of the family welfare function and individuals' utilities, and the "technologies" for home production.

# Table 1. Pension Eligibility and Pension Receipt

Percentage of individuals in each age group reporting pension receipt

| | Non-Blacks | | Blacks | |
|---|---|---|---|---|
| Age Group | Males | Females | Males | Females |
| 50-54 | 2 | 1 | 2 | 2 |
| 55-59 | 4 | 6 | 3 | 8 |
| **60-64** | 17 | **40** | 15 | **60** |
| **65-69** | **36** | 55 | **61** | 85 |
| 70-74 | 39 | 61 | 82 | 89 |
| 75-79 | 40 | 55 | 84 | 90 |
| 80-84 | 36 | 56 | 93 | 93 |
| 85+ | 51 | 56 | 89 | 90 |

Notes: Source is Income and Expenditure Survey, 2000. Cells are weighted according to the survey's sampling probabilities, so that indicated percentages represent population level estimates.

Table 2: Budget shares and pensions

Instrumental variables
Pension recipient(s) in HH
(Instrument: Age elig indiv in hh)

| *Budget share* | Male Only (4) | Female Only (5) | Both M&F (6) |
|---|---|---|---|
| *1. Food* | -3.4<br>(4.1) | 6.3<br>(1.8) | -3.5<br>(4.4) |
| *2. Human capital* | | | |
| 2.1 Health care | 1.4<br>(0.8) | 0.8<br>(0.8) | 2.4<br>(1.1) |
| 2.2 Schooling | 0.17<br>(0.82) | -0.48<br>(0.79) | 0.12<br>(1.2) |
| *3. Child goods* | | | |
| 3.1 Baby food | 0.30<br>(0.33) | 0.193<br>(0.183) | 0.650<br>(0.387) |
| 3.2 Girls' clothing | 0.82<br>(0.40) | -0.23<br>(0.33) | 0.72<br>(0.54) |
| 3.3 Boys' clothing | -0.03<br>(0.99) | -0.17<br>(0.28) | 0.086<br>(1.01) |
| *4. Adult goods* | | | |
| 4.1 Adult female clothing | -0.65<br>(0.69) | 0.68<br>(0.50) | 0.66<br>(0.87) |
| 4.2 Adult male clothing | 0.56<br>(0.58) | -0.91<br>(0.38) | -0.47<br>(0.69) |
| 4.3 Alcohol | 0.92<br>(1.3) | -0.010<br>(0.54) | 1.1<br>(1.4) |
| 4.4 Tobacco | 0.10<br>(0.84) | -0.76<br>(0.36) | 0.13<br>(0.90) |

**Notes**: Income and Expenditure Survey, 2000. Sample size is 20,643 households. Standard errors are computed using "sandwich" estimator of variance/covariance matrix. Regressions also include controls for lnPCE, number of males and number of females in household, broken down into 5-year age groups and magisterial district fixed effects.

Table 3. Living Arrangements and Pensions

| Dependent Variables: Household composition | | Explanatory Variables: Gender of Pension Eligible Individual | | | | | |
|---|---|---|---|---|---|---|---|
| | | Sole Male | | Sole Female | | Both Male and Female | |
| Females | Age 0-9 | 0.15 | (0.21) | 0.16 | (0.066) | 0.33 | (0.22) |
| | Age 10-19 | -0.007 | (0.23) | 0.079 | (0.070) | 0.075 | (0.24) |
| | Age 20-29 | 0.088 | (0.18) | -0.14 | (0.056) | -0.035 | (0.19) |
| | Age 30-39 | -0.14 | (0.14) | -0.19 | (0.042) | -0.34 | (0.14) |
| | Age 40-49 | -0.10 | (0.11) | -0.11 | (0.035) | -0.24 | (0.12) |
| Males | Age 0-9 | 0.065 | (0.21) | 0.088 | (0.065) | 0.21 | (0.22) |
| | Age 10-19 | -0.11 | (0.23) | 0.24 | (0.071) | 0.038 | (0.24) |
| | Age 20-29 | -0.10 | (0.19) | 0.053 | (0.058) | -0.021 | (0.20) |
| | Age 30-39 | -0.009 | (0.14) | -0.000 | (0.044) | 0.025 | (0.15) |
| | Age 40-49 | 0.014 | (0.12) | -0.062 | (0.036) | 0.036 | (0.12) |

Notes: Source is Income and Expenditure Survey, 2000. Sample size is 20,643 households. Standard errors (computed with the sandwich estimator of the variance/covariance matrix) are reported in parentheses beside coefficient estimates. Regressions were estimated by seemingly unrelated regression. The 10 equations estimated correspond to the 10 dependent variables listed in the left column. Each column represents the coefficient estimates for pension eligibility variables. Also included in each regression was the number of males in females in all other age groups.

Table 4: Adult education, adult height and pensions
Instrumental variables estimates (Instrument: Pension Eligible in HH)

| | | Education | | Height | |
|---|---|---|---|---|---|
| | | Male | Female | Male | Female |
| | | (1) | (2) | (1) | (2) |
| *1. Any pension recipient in household* | | | | | |
| Age 20-55 | | -0.74 | -0.50 | 2.8 | -4.0 |
| | | (0.61) | (0.46) | (1.8) | (1.6) |
| Age 20-30 | | -0.25 | -0.74 | 3.4 | -4.0 |
| | | (0.85) | (0.59) | (3.0) | (2.3) |
| Age 31-55 | | -1.1 | -0.22 | 2.0 | -4.7 |
| | | (0.73) | (0.67) | (2.2) | (2.1) |
| *2. Separate only male recipients from female or both male and female recipients in household* | | | | | |
| Age 20-55 | | | | | |
| (1) if only male recipient | | -1.7 | -0.73 | 4.6 | -3.9 |
| | | (0.77) | (0.58) | (2.4) | (1.8) |
| (1) if female/both M&F recipients | | 1.4 | -0.008 | -0.28 | -4.3 |
| | | (0.92) | (0.61) | (2.4) | (2.8) |
| Difference (F/Both - M) | | 3.1 | 0.72 | -4.8 | -0.41 |
| | | (1.2) | (0.79) | (3.2) | (3.2) |
| Age 20-30 | | | | | |
| (1) if only male recipient | | -1.5 | -1.4 | 6.2 | -2.6 |
| | | (1.2) | (0.82) | (4.5) | (2.8) |
| (1) if female/both M&F recipients | | 1.9 | 0.53 | 0.13 | -8.2 |
| | | (1.1) | (0.58) | (3.4) | (4.3) |
| Difference diff (F/Both - M) | | 3.4 | 1.9 | -6.0 | -5.6 |
| | | (1.7) | (0.97) | (5.2) | (4.9) |
| Age 31-55 | | | | | |
| (1) if only male recipient | | -1.9 | -0.22 | 2.2 | -5.5 |
| | | (0.82) | (0.81) | (2.5) | (2.2) |
| (1) if female/both M&F recipients | | 0.91 | -0.22 | 1.2 | -1.7 |
| | | (1.2) | (0.96) | (3.4) | (4.5) |
| Difference (F/Both - M) | | 2.8 | 0.004 | -1.1 | 3.8 |
| | | (1.4) | (1.2) | (3.9) | (4.8) |
| Sample sizes: | All | 9,363 | 11,437 | 3,509 | 4,743 |
| | 20-30 | 4,009 | 4,407 | 1,444 | 1,856 |
| | 31-55 | 5,354 | 7,030 | 2,065 | 2,887 |

Notes: Source is 1998 Demographic and Health Survey. Not all adult heights were measured. Regressions also control respondent's age and education, household composition and wealth. Instruments are eligibility for pension. Robust standard errors in parentheses.

Table 5: Child educational attainment and pensions

| Specification | Males | Females |
|---|---|---|
| | (1) | (2) |
| *1. Intent-to-Treat* | | |
| *Any age-eligible person in HH* | -0.11 | 0.017 |
| | (0.11) | (0.10) |
| *2. Instrumental variables* | | |
| *A. Any pension recipient in HH* | | |
| All children | -0.19 | 0.029 |
| | (0.22) | (0.21) |
| Young children | -0.30 | 0.37 |
| | (0.20) | (0.21) |
| Older children | -0.11 | -0.40 |
| | (0.39) | (0.37) |
| *B. Female recipients relative to male and couple recipients* | | |
| All children | | |
| Male/both male and female recipients | 1.21 | 0.42 |
| | (1.22) | (0.76) |
| Only female recipients | -0.11 | 0.045 |
| | (0.24) | (0.21) |
| Difference (Male/Both - Female) | 1.3 | 0.38 |
| | (1.1) | (0.72) |
| Young children | | |
| Male/both male and female recipients | -0.12 | 1.5 |
| | (1.1) | (0.81) |
| Only female recipients | -0.30 | 0.35 |
| | (0.21) | (0.24) |
| Difference (Female - Male/Both) | 0.18 | 1.1 |
| | (0.99) | (0.84) |
| Older children | | |
| Male/both male and female recipients | 2.3 | -2.0 |
| | (1.4) | (0.85) |
| Only female recipients | 0.069 | -0.60 |
| | (0.40) | (0.39) |
| Difference (Female - Male/Both) | 2.3 | -1.4 |
| | (1.2) | (0.68) |

Notes: Source is Demographic and Health Survey, 1998. There are 9,422 male children age 6-19, of whom 5,044 are younger males (age 6-12) and 4,378 are older males (age 13-19); there are 9,206 female children of whom 5,058 are younger and 4,148 are older.

Table 6. Differential Apparent Treatment Effects, by Mother's Coresidence

| Years of Education Attained | Boys | | | Girls | | |
|---|---|---|---|---|---|---|
| | Ages 6-15 | Ages 6-10 | Ages 11-15 | Ages 6-15 | Ages 6-10 | Ages 11-15 |
| Coresident with pensioner | -0.18 | -0.33 | -0.015 | 0.072 | 0.34 | -0.30 |
| | (0.25) | (0.23) | (0.41) | (0.25) | (0.22) | (0.45) |
| Coresident with mother | 0.16 | 0.094 | 0.28 | 0.21 | 0.14 | 0.31 |
| | (0.076) | (0.076) | (0.13) | (0.071) | (0.072) | (0.13) |
| Coresident with pensioner *and* mother | -0.42 | -0.15 | -0.75 | -0.23 | -0.17 | -0.30 |
| | (0.14) | (0.14) | (0.25) | (0.13) | (0.13) | (0.22) |

Notes: Source is Demographic and Health Survey, 1998. There are 6657 boys aged 6-15, of whom 3643 are aged 6-10 and 3014 are aged 11-15. There are 6800 girls, of whom 3628 are aged 6-10, and 3172 are 11-15. The age range corresponds with the age range for which DHS reports mother's coresidence. As in previous regressions, age eligibility is used to instrument for pension receipt, and child's age (in single year dummies) and a vector of household demographic controls are also included in the regression.

Table 7. Children's Educational Attainment and Present and Future Pension Income

| Years of Education attained by the year 1993 | Boys and Girls Ages 6-19 | Girls | | Boys | |
|---|---|---|---|---|---|
| | | Ages 6-12 | Ages 13-19 | Ages 6-12 | Ages 13-19 |
| Coresident with pension eligible individual in 1993 | 0.09 | 0.039 | 0.17 | 0.22 | 0.003 |
| | (0.10) | (0.11) | (0.25) | (0.085) | (0.27) |
| Coresident with pension eligible individual in 1998 | 0.24 | 0.28 | -0.091 | 0.20 | 0.11 |
| | (0.13) | (0.12) | (0.29) | (0.14) | (0.36) |
| Coresident with eligible in 1993 *and* 1998 | -0.37 | -0.31 | -0.16 | -0.39 | -0.34 |
| | (0.17) | (0.17) | (0.41) | (0.17) | (0.48) |

Notes: Source is the KwaZulu Natal Income Dynamics Survey (KIDS), 1998. The age categories are defined according to the child's age in the year 1993. Also included in the regression- Age of child (dummy for each single year of age), numbers of elderly males and females resident in the household in 5 year age bands. There are 2554 children, of whom 672 are younger girls, 621 are older girls, 707 are younger boys, and 554 are younger girls.

Figure 1. Pension Uptake Around the Age of Eligibility