

# **Pensions and Priors: A South African Case Study**

Seán M. Muller

[This version: 3 December 2007]

# 1 Introduction

Rather than merely abandon the poor and economically weak to the vicissitudes of the economy and markets, most modern day governments utilise – to a greater or lesser extent – a variety of policies to ease the plight of these individuals, improve their prospects and increase their economic power. However, many such policies are the subject of intense criticism from certain factions in society, in which economists are often disproportionately represented. The rationale for these critiques is, as noted by Galbraith (1987: 123), often related to the idea that aid only serves to continue the impoverishment of the poor. This is particularly problematic since much of this criticism appears to be more a function of particular preferences for some theoretical priors over others, even in the face of contrary empirical evidence.<sup>1</sup> One such example in the South African context relates to the provision of a substantial social pension.

Despite work finding significant positive impacts of the social pension, by Case and Deaton (1998), and Duflo (2003), the Department of Social Development included the possible negative externalities of this policy as part of a recent investigation into “allegations of perverse incentives in the government's social grant system”.<sup>2,3</sup> This followed the conclusion of Bertrand, Mullainathan and Miller (2003) that “African household members 16-50years old reduce their labour supply when they live with pension beneficiaries”, and hence that “policy-makers in developing countries would do well to consider how different living arrangements could interfere with their social objectives” (Bertrand, Mullainathan and Miller: 48-49).<sup>4</sup>

Ranchod (2006a) using a somewhat different approach – discussed further in section 2 - comes to similar conclusions. Ranchod (2006b) finds that the pension reduces labour supply amongst the elderly, leading him to suggest that “the state needs to carefully consider the incentives it provides under its various welfare programmes” (2006b: 744).<sup>5</sup>

Ostensibly, the primary focus of the work by Bertrand et al. (henceforth BMM) was not labour supply per se, but rather the distribution of social grants within a household; the concern being that if a grant is intended for one particular household member, its

---

<sup>1</sup> Leonard (2000: 140) for instance has noted, in the context of the debate over the employment and welfare effects of the minimum wage, the disciplinary tension in economics that manifests itself when “compelling evidence meets core theoretical belief”.

<sup>2</sup> Social Cluster Briefing, 11 May 2006

<sup>3</sup> Case and Deaton found that the pension does reach particularly poor households and is used in much the same way as other income (i.e. is not frittered away or spent particularly irresponsibly), whilst Duflo found positive impacts on children's health status.

<sup>4</sup> Work by Ranchod (2006a and 2006b) has come to similar conclusions regarding labour supply. Those studies are discussed further in section 2 below.

<sup>5</sup> This inference is arguably flawed. First, it is not clear how one could differentiate between a norm regarding retirement and such a result – though of course some would no doubt argue that a norm would not evolve unless it were based on an underlying tendency to favour labour withdrawal in such a situation. Perhaps as important, it is far from clear that such a reduction in labour supply could be considered a negative social outcome.

intended effects may be subverted by ‘capture’ or voluntary redistribution within the household. Because the distribution of income and resources is rarely observable, they focus instead on the labour supply of prime-age individuals as an indication of the extent to which a given individual benefits.<sup>6</sup>

In the working paper version BMM (2000: 5-8) outline a standard individual choice model based on the trade-off between leisure and labour. They note that pension receipt may affect labour supply through two channels: Through an ‘income effect’; and through an ‘incentive effect’. The former occurs as a lump-sum transfer and therefore is hypothesised to increase leisure time because this is a normal good. The latter follows from what BMM suggest may be pensioners’ desire to insure other members of the household against shocks. As a result of that motivation, prime age individuals’ receipt of pension income is negatively related to their income from employment. Hence the disincentive effect on labour supply.

This paper begins by assuming a different theoretical prior to that of BMM. That is, that *household formation*, rather than the labour supply decision, is the process of primary importance in examining the impact of the social pension. The hypothesis that followed from BMM’s prior was that pension receipt would induce labour supply withdrawal. The hypothesis that follows from the household formation perspective is that pension receipt may induce the presence of individuals who are less likely to be actively participating in the labour force.

Section 2 provides a brief outline of the BMM rationale and argument, critiquing aspects in some detail and noting additional implications of that hypothesis not addressed in their paper. Section 3 provides a basic model of household formation based on income-smoothing across households of an extended family, and notes some of the key predictions that follow from it. Section 4 outlines the methodology of the investigation. Section 5, the main empirical section, presents the results. Section 6 discusses the findings and possible interpretations thereof. Section 7 concludes.

The primary aim of this paper is demonstrating that, based on existing evidence, one can make *at least as strong* a case that the pension influences household formation, as that it affects labour supply. Whilst we are of course concerned with determining the true impact, the nature of data used prohibits the drawing of definitive conclusions

---

<sup>6</sup> Analogously, Duflo(2003) examines differences in children’s weight-for-height and height-for-age associated with pension eligibility in order to determine the extent to which children benefit from pension receipt.

## 2 Intra-Household Redistribution and Labour Supply

As noted above, BMM postulate two channels through which the social pension may (negatively) affect labour supply; an ‘income’ and an ‘incentive’ effect. The extent to which these two effects have an impact depends on the issue upon which the analysis began; namely, the distribution of (grant) income within the household. BMM note two key factors determining distribution: *bargaining power* and *altruism*. The two are theoretically, and particularly empirically, difficult to disentangle. Furthermore, there are a number of factors relating to pension recipients and household members that could impact on these variables. Amongst them are age, gender and familial relation, to name the most notable. As an example, the evidence from Duflo (2003) suggests that female pensioners may be more likely than their male counterparts to distribute resources toward female household members. BMM’s primary hypothesis in this regard, though based largely on intuition, is that male household members are likely to receive/appropriate more than their ‘fair share’.<sup>7</sup> Their findings support this, in as much as pension receipt is found to be significantly associated with a lower labour supply for men but not women.

### IDENTIFICATION

The primary identification strategy used by BMM hinges around comparing the employment status of working-age/‘prime age’ individuals in pension eligible households with those in non-eligible households. Using the pension amount received as the main regressor of interest in a regression predicting employment status is one obvious strategy. However, this may well be biased by endogeneity in the take-up of pensions: in other words, there may be something structurally different about households that take advantage of their eligibility relative to those that do not. If those structural differences affect employment levels or employment status then the results will be biased. The preferred specification therefore is one in which total pension income received by a household is instrumented by the number of pensioners in that household.

Whilst using an IV eliminates one form of bias, there remain other possible sources. As BMM note (2000: 13), the main concern is that there are systematic differences between individuals in pension-eligible and pension-ineligible households. They propose to deal with that problem using a three-step procedure: By restricting the analysis to three-generation (henceforth 3G) households; showing *within* this subset that the impact on labour supply is associated specifically with pension eligibility rather than merely the presence of elderly members (i.e. demonstrating an age discontinuity); and, looking for any systematic differences on observable characteristics between prime age individuals in the two types of 3G households.

BMM also note two ‘additional possible sources’ of bias: that the elderly may induce different labour market behaviour because they require care; and that there may be selective migration resulting from the pension. They propose addressing these by

---

<sup>7</sup> Where ‘fair’ here refers simply to equal distribution amongst (at least prime age) household members.

controlling for the health of the elderly, and ascertaining whether pension income has a positive impact on being a migrant within a household (using a variable which records whether an individual migrated into the household in the previous 5 years).

In short, BMM's primary focus is on: the difference in employment levels and status of black African males, of working age, in three generation households that are pension eligible and ineligible.<sup>8</sup>

In addition to that study, Ranchod (2006a) attempts to address the same issue using panel data – which is, as all authors have noted, the most appropriate for answering these questions. Specifically, he uses a panel from September 2001 to September 2003 constructed from the Labour Force Surveys. Rather than directly examining the impact of pension receipt, he looks at the inverse issue: the impact of pensioner *death or out-migration* on labour supply. His finding is that there is evidence of increased labour supply as a result of these 'shocks'; again suggesting that higher unemployment/non-participation in households with pensioners may be due in large part to a voluntary reduction in labour supply following from pension receipt.

## FURTHER ANALYSIS AND EXTENSIONS

As noted, the studies by BMM and Ranchod both find a statistically significant association between pension receipt and lower levels of labour force participation amongst household members. In addition, this result is given a causal interpretation – that prime age African household members *reduce* their labour supply when they live with pension beneficiaries. Posel, Fairburn and Lund (2004), question the breadth of this conclusion. They show that when a more inclusive definition of the household is used – i.e. including non-resident members – pension receipt may be associated with higher labour force participation in some instances (specifically in the case of *female participation* and *female pension receipt*). Underlying that critique, and the analysis here, is a scepticism of the BMM result based upon the view – also expressed by Case and Deaton (1998: 1331) – that one would expect the distortionary effect on labour supply of a cash transfer like the pension to be insignificant given the South African context of poverty and high unemployment.

Whilst we do replicate here one possibly important result noted in Posel et al's paper – relating to the econometric impact of different definitions of 'household size' - the emphasis is on assessing the validity of the BMM approach for the dependent variables on which they focused: the labour supply of *resident* household members.

Exploiting the non-linear effects at the age of pension eligibility BMM do indeed, as they claim (Bertrand et al., 2003: 41), quite convincingly demonstrate that the identified effect is that of the pension and not simply the effect of living in a household with elderly

---

<sup>8</sup> The focus on black African individuals arises because they constitute the vast majority of recipients, and are more likely to comprise 3G households.

occupants. However, an issue not fully interrogated in the BMM rationale is that which relates to the *primary policy concern*. Are we concerned that income received by other household members is *generally* subject to capture and/or redistribution, and hence has a negative effect on labour supply? Or that *only* pension income plays such a role – because it is more prone to capture, or pensioners are more altruistic? Whilst the latter issue is raised by BMM, the implications of this rationale being extended to broader household income are not discussed.

For policy purposes, however, they are important. If pension income has the same effect as other household income it is not clear that any consequences thereof can be directly attributed to the pension *per se*. One cannot then argue that it is the *pension* that has a distorting effect on labour supply, but only that increased household income generally has this effect. Thus the foremost concern from a policy perspective is with *differential* impacts of social grants. In this respect, evidence from experimental economics – and in particular, ‘dictator games’, see Cherry (2001) – suggests that unearned income is more likely to be redistributed (voluntarily or otherwise) than earned income. This does bolster the case for focusing on pension impacts as likely to be different from those of other forms of household income.

In addition to these broader concerns, there are two further implications of the BMM hypothesis that are important.

First, if pension income is associated with labour supply reduction, this should also be visible not only in the binary choice over employment status, or the (possibly) continuous variable of hours worked, but also in the decision to search or not search conditional on being unemployed. In other words, we can also ask whether there is a differential between the average labour market status of *non income-earning individuals* in pension eligible, and pension ineligible households.

The reasoning is as follows. BMM suggest that labour supply of those employed individuals resident within the household is reduced in response to pension receipt because of an incentive – a share of the pension – associated with unemployment. That being so, an *unemployed, searching* individual in an eligible household also has the same incentive not to engage in job *search*. This is because the incentive to search is implicitly reduced by the availability of pension income conditional on being unemployed. One can make this argument more formally using a standard search model of the kind presented by Mortensen and Pissarides (1999: 2571):

$$U_t = \frac{b - a}{1 + r} + \frac{1}{1 + r} \int \max \{ W, U_{t+1} \} dF(W),$$

$$t = 1, 2, \dots$$

In this Bellman equation,  $r$  is the discount rate,  $b$  is income flow received contingent on unemployment, and  $a$  represents the cost of search per period,  $W$  is the current value of the highest available wage offer (if there is one),  $F(W)$  is a cdf characterising the distribution of payoffs and  $U_{t+1}$  is the value of continued search. For our purposes,  $b$  is the primary variable of interest: pension receipt and distribution thereof based on unemployment would increase the value of  $b$ , thereby reducing the incentive to search given that an individual is already unemployed.<sup>9</sup>

The second implication is that, *ceteris paribus*, one should expect that as the amount of pension income received by a pension-receiving household increases, so should the negative effect on labour supply also increase. This follows from the larger amount of income available for distribution/appropriation.<sup>10</sup>

These two predictions are tested in section 4 and discussed in section 5 below.

One final matter worth noting in the BMM analysis is their approach to unemployment. They suggest that those who are involuntarily unemployed should be substantially different – whether on observed or unobserved characteristics - from those who are employed. As a result, if it is true that there is some kind of endogenous household formation going on whereby unemployed individuals select into pension-eligible households, then we should expect to see differences in variables that typically affect the probability of finding employment. As one test of this, BMM analyse education levels across prime age individuals in eligible and non-eligible households, find no significant differences and conclude that household formation of this sort is unlikely to be a problem. However, this logic is problematic. In a high unemployment, low skills environment, there is every reason to believe that to a significant extent being (un)employed is a matter of (bad) luck. This is especially true at the very low-skilled end of the labour market which is what is predominantly captured in the BMM analysis. Thus, those selecting (‘positively’ or ‘negatively’ – see Section 3) into pension-eligible households are not doing so necessarily because they have *structurally* different prospects of finding employment, but simply because they got the short-straw in the employment draw. If this is an accurate assessment of the situation – as it would seem to be - then we cannot infer what BMM suggest about household formation from (the lack of) differences in educational attainment across eligible and ineligible households.

---

<sup>9</sup> One might argue that the extent of the intra-household transfer might not be large enough to induce a detectable change of this sort – but that is incompatible with BMM’s proposal that it *is* sufficiently large to induce voluntary withdrawal from *existing employment*.

<sup>10</sup> Clearly this relies on, amongst others, the assumption that there is no gross limit to the proportion of pension income that is received/can be appropriated by an unemployed, prime-age household member. For instance, if there were some implicit rule such that unemployed individuals in the household do not receive more than R50/month in intra-household transfers, we would not expect to see labour supply decreasing concomitantly with pension income.

### 3 Inter-Household Redistribution and Household Formation

In addition to the theoretical model described above and its application to the PSLSD data, BMM (2000 and 2003) briefly discuss some qualitative data on the effect of the social pension which they suggest supports the hypothesis that pension receipt induces a reduction in labour supply:

*Anecdotal evidence and newspaper articles hint that the pension may well have affected relatives' labour supply. One article mentions that "the impact of pensions in communities with a high rate of unemployment was huge, as multi-generation households formed a constellation around the person receiving the pension"... Another describes a pensioner's "five children, who also live with him in his two-bedroom flat, contribute to the family income when they can find work. But none has a full-time job".*

(Bertrand et al., 2003: 28).

The first of these anecdotes seems to quite clearly suggest household formation in response to pension receipt. However, although BMM rule-out the possibility of such a process having any significant effect on their analysis, it may nevertheless be consistent with the broader hypothesis of a reduction in labour supply. In particular, working individuals leaving their employment in order to migrate into a pension-eligible household would still be consistent with the first story above. The second anecdote mitigates against that interpretation though, as it suggests that (in that instance at least) the pensioner's prime age dependents are prepared to work but unable to find employment.

Interpretation of such evidence is necessarily tentative and inconclusive; nevertheless it would appear that these anecdotes are in fact more supportive of a process of selection into pension-eligible households by those struggling to find employment, rather than the voluntary withdrawal of individuals from the labour force.

#### MODELLING HOUSEHOLD FORMATION

The nature, processes and causes of migration into and out of households are far from fully understood. Thus, in order to clarify the proposed alternative to the BMM hypothesis, we construct below a simple model which captures the primary intuition behind proposing that household formation is induced by pension receipt.

We begin with an extended family consisting of  $N$  members, separated into two households, which maximises, implicitly or explicitly, a family utility function  $U^F$ .<sup>11</sup>

---

<sup>11</sup> This may be the result of 'fairness' norms of the kind derived by Rawls (1971) – but applied at the household level - and/or one could give it some basis in genetics. Following Binmore (1993 and 1998), one would expect that if such a norm does exist it would have some foundation in power relations and/or genetic interest.

$$(1) \quad U^F = \sum_{i=1}^n u_i(y_i)$$

$$\text{where } \partial u_i / \partial y_i > 0, \partial^2 u_i / \partial y_i^2 < 0$$

$y_i$  is the amount of income allocated to individual  $i$ ,  $n$  is the total number of family members,  $u_i$  is  $i$ 's utility function and it is assumed that  $u_i(x) = u_j(x)$  for all  $i, j$  and  $x$ .<sup>12</sup>

In the simplest case in which there is one individual, either a pensioner (P) or a worker (W), receiving income in the household, the manner in – and indeed extent to – which that income is distributed is assumed to be a function of that individual's broader utility function (' $v$ ') as follows:<sup>13</sup>

$$(2) \quad V_P = u_P(y_P) + \alpha_P \sum_{i=1}^{n_P-1} u_i(y_i)$$

$$V_W = u_W(y_W) + \alpha_W \sum_{j=1}^{n_W-1} u_j(y_j)$$

$$j \neq i, 0 < \alpha < 1, \alpha^W \leq \alpha^P$$

Alpha measures the weight the *income recipient*, either pensioner (P) or worker (W), places upon the welfare of the other household members. It is assumed that pensioners are more generous in this sense than workers.<sup>14</sup>  $n_z$  is the number of household members in the respective household,  $z \in \{p, w\}$ .

Non-earning individuals maximise their welfare by moving to households where the proportion allocated to them is the greatest. In the case where  $\alpha^P = \alpha^W = 1$ , and there is free entry to the household with zero 'transaction costs' (i.e. no cost to moving) this will serve to maximise (1) given (2) with  $y_i^h = (Y^h/n)$ . Where  $\alpha^W < \alpha^P \neq 1$ , it will be equivalent to maximising (1) under the constraint that for all  $h$ , and  $z$ :

$$(3) \quad \alpha_z (\partial u_i^h / \partial y_i^h) = \partial u_z^h / \partial y_i^h$$

<sup>12</sup> Note that since we do not sum over 'h', this effectively assumes that individuals only receive money via the household, so that an individual never receives money directly from two households. They can however receive money from another household *via* remittance and distribution within their household of residence – see equation (4).

<sup>13</sup> As it stands this is potentially consistent with both unitary and collective models of such a hypothetical household.

<sup>14</sup> As noted in the previous section, there is some support for this hypothesis; not least that a social pension may in some sense be unearned and hence a pensioner would feel less entitled to the income than a worker would to their wage. This discrepancy between distribution of ostensibly earned and unearned income has been found in various experiments – see for instance Cherry (2001).

A more general model would allow for the possibility that individuals within the family also care about family members *outside* of their own household. This can be achieved by removing the household superscript in the income earner's utility function so that they will maximise their utility by distributing income amongst all family members (N) such that they care also about family members in other households:

$$(4) \quad \alpha_z(\partial u_i / \partial y_i) = \partial u_z / \partial y_i$$

This is consistent with empirical evidence of significant remittances between households (in addition to those sent by an 'absent'/migrant household member). The extent of remittances in general can be seen from the 2003 General Household Survey in which approximately 14% of households reported remittances as being their *primary* income source.

A further generalisation is to allow for the possibility of different individuals having different weightings in the decision-maker's utility function. This is necessary to explain any differences in effects of marginal changes in income between individuals within the household. To align this extended model with 'collective' models of the household, these coefficients could be made to represent the outcome of an intra-household bargaining process.<sup>15,16</sup>

## TYPES OF FORMATION

Given the model above, two key questions arise:

1. What behaviours constitute 'household formation'?
2. If individuals can send income to other households why would we observe physical movements at all?

In terms of household formation, we will be primarily concerned here with distinguishing between two types: 'positive' and 'negative'. Positive formation occurs when an individual moves into (out of) the household type of interest due to a structural difference between that household and the individual's previous (new) residence.<sup>17</sup> Negative formation occurs when an individual *stays in* (does *not* move into) the household type of interest as a result of some structural difference between it and the alternative (current) residence.

Some examples may assist in clarifying these assertions. In the case of pension receipt: individuals' moving into pension-receiving households rather than non-receiving households due to the additional income available, constitutes positive formation. Unemployed children's remaining in a 3G household because it receives a pension,

---

<sup>15</sup> Alderman et al. (1995) argue that such models should be the baseline rather than the exception.

<sup>16</sup> Note, however, that the fact that individuals receive different weightings is not necessarily inconsistent with the unitary model.

<sup>17</sup> Although random movements might be of interest for some economic issues they are not salient here.

whereas they *would have moved out* in the (counterfactual) event that they received a job or the household was without a pension, constitutes negative formation.

The model above, whilst explicitly phrased in terms of positive formation does implicitly capture the (inverse) rationale behind negative formation. Note that for the purposes of the analysis that follows, identification of the ‘reference household’ is particularly important. Following BMM, these will be three-generation households – either pension eligible or not. Thus pension receipt is the variable that induces a structural difference between eligible 3G households and individuals’ unspecified outside alternative that is not present for non-eligible 3G households. Hence the hypothesis that the former are subject to either positive formation via individuals migrating into the household in order to benefit from the additional funds, or negative formation whereby individuals remain within the household for the same reason.

A third type of formation, that may also explain – in different ways - the BMM and Ranchod (2006a) results, is the *dissolution* or *non-dissolution* of a household. This can be in terms of the household’s type – continuing as a 3G household or not – or its actual existence. For the BMM results that focus on 3G households: If it were the case that 3G households *as a type* only continue to exist if they: i. receive a pension, or ii. receive a certain amount of income from employment, then this could be sufficient to generate the finding that pension-receiving 3G households have fewer employed members (those 3G households that do not receive a pension dissolve unless they have more working members). For the Ranchod result, because of his focus on *pensioner loss*, the logic reverses and is based on *existence*: if households (the study is not limited to 3G households) lose a pensioner they will only stay together if they are able to find additional employment. In that case, the only households in existence after pensioner out-migration/death will indeed have higher employment levels.<sup>18,19</sup>

The issue of migration versus remittances is obviously critical to determining the implications of family income-smoothing across households. If smoothing primarily takes place through remittances we are unlikely to pick-up its consequences by analysing household structure, composition and characteristics – instead more detailed knowledge of remittance behaviour (especially matching sending and receiving households) would be required. However, there are good theoretical reasons for believing that smoothing would take place via physical movement. By having the needy individual within the household, the primary income-earner can directly control the extent of resources

---

<sup>18</sup> Notice that even panel-level analyses such as Ranchod’s can therefore suffer from biases, where these involve counterfactuals and/or issues of existence. A particular problem for that analysis is the fact that pensioner out-migration – which accounts for most ‘pensioner loss’ – is inherently endogenous.

<sup>19</sup> One could possibly control for this to some degree by looking at the overall attrition rate of households in the panel and ascertaining which ones were initially pension-eligible.

allocated to that individual, and also monitor their behaviour (and therefore make an informed determination of ‘need’) to a much greater extent.<sup>20</sup>

## IMPLICATIONS

For the purpose of empirically assessing the above model, and especially doing so relative to the one of BMM, it is critically important that its implications be clearly outlined.

To start with, the model suggests that in estimating the impact of the pension we should control for household income. This is because the smoothing migration that occurs is a function of the income an individual believes they are likely to receive if they migrate into, or remain in, a household. Assuming the base scenario with alpha equal to 1, it would seem most appropriate to control for household per capita income.<sup>21,22</sup> Note at this point however that the use of this control is not likely to affect an analysis of the kind conducted below. The reason is that it is based on the PSLSD, a static, cross-sectional dataset, and hence captures data such as household size that will *already* incorporate any migration that has taken place. Thus we anticipate similar per capita incomes across eligible and ineligible households.<sup>23</sup>

To the extent that positive migration is important – i.e. there are prime age individuals who can choose between households with pension-eligible individuals and those without (3G or otherwise) – the model predicts that there should be a larger number/proportion of migrants in *eligible* 3G households. This can be tested to a degree, using the question in the PSLSD survey which asks whether a given household member migrated into the household in the five years prior to the survey. Though there is the possibility that the five year time span may be rather limited in terms of capturing the full impact of this behaviour.<sup>24</sup>

---

<sup>20</sup> Although we have assumed zero transaction costs, an additional argument for physical income smoothing across households is that a once-off migration cost is likely to be cheaper than the net present cost of future, regular financial remittances.

<sup>21</sup> In doing so it is necessary – at least in interpreting the significance of coefficients - to be aware of possibly severe multicollinearity between household per capita income and household size.

<sup>22</sup> Though in their analysis BMM assume the opposite extreme, that alpha equals 1 for prime-age males (in terms of distribution of pension income) and 0 for all other household members.

<sup>23</sup> Recall that we are assuming that possible dependents – i.e. members of the extended family – of both types of household are similar in their outside (non-3G household) options. The assumption of the household formation model outlined above was that migration equalises per capita income across alternatives. And in order for that to happen, pension-receiving households must, *ceteris paribus*, be the subject of greater positive and negative migration.

<sup>24</sup> By contrast, if the oft-made assumption that the extent of the pension for Africans was marginal prior to the 1990s is true, then one might think that the variable is sufficient to capture most of the migration that pension receipt would have induced prior to the PSLSD. Case and Deaton (1998: 1334) in fact suggest that the PSLSD survey was too soon after the pension program moving into ‘full operation’ to capture migration effects.

It is less clear how we might assess the presence of ‘negative migration’. Aside from panel data, a variable that captured the original number of children in a household might allow for some assessment of child out-migration. However, no such information is present in the PSLSD, and recall that our intention here is to assess the validity of the BMM conclusions relative to a household formation hypothesis *using the same data source*. Even with panel data, such an approach would require a much more detailed knowledge of the actual household structures involved, over and above the fact that they are 3G households of some sort. In addition, capturing the second type of negative formation – in which individuals do not move into the household of interest – is virtually impossible.

Due to these types of migration, the household formation model (henceforth HF model) predicts that pension receipt will be associated with lower hours worked, and a lower probability of employment amongst those who are likely to be able to receive some of the pension income if they migrate (i.e. if the pensioners are prepared to ‘insure’ them against the shock they have experienced). So if this model is true, then the BMM estimation procedure is inherently biased. The key difference in terms of policy from the labour supply model (henceforth LS model), and the main emphasis of this paper, is in the causal interpretation that follows from observing the above-mentioned differences. It must be emphasised that this interpretation follows from the fact that the two causal mechanisms are *fundamentally* different: the LS model proposes that lower labour supply in pension-eligible households results from a reduction in labour supply by resident members; the HF model proposes that it is merely the manifestation of *already* unemployed or under-employed individuals entering (or remaining in) such households.

The fact that the two models make similar predictions in this regard is of course not very helpful for assessing their relative merits. However, in relation to the issue of employment status there is one key respect in which the two differ: As noted in Section 2, the LS model also implies that, *conditional on being non-employed*, prime age individuals will be less likely to engage in job search if they are resident in a pension-eligible household. This is because pension income is hypothesised to reduce labour supply. If that is the case, it should manifest not only in the probability of being employed but also in the conditional probability of being engaged in job search. By contrast, the HF model gives us no reason to believe that individuals who are non-employed in pension-receiving and non-pension receiving households will be any different in their propensity to engage in search.<sup>25</sup>

This difference in model implications is important for another reason; it potentially allows us to address the existence problem noted on p.10 whereby ineligible households inherently have more employed individuals simply because they would not *exist*

---

<sup>25</sup> One might argue that the effect of the pension on assisting job search may mask the impact of this on labour supply. However, such an argument effectively makes the point; from a policy perspective the primary concern - particularly in light of the worries about ‘perverse outcomes’ - is with *aggregate impacts*. I.e. If the negative effect of the pension on search for some is offset by its enabling search by others then it cannot be said to be having deleterious effects on the labour market.

otherwise. If there is per capita income smoothing over households (see footnote 22) comparing average labour supply per working-age household member across these households will always find a higher value in those households where there is a wage earner (who *by definition* is supplying a substantial amount of labour).<sup>26</sup>

The main point then that follows from the HF model is that the BMM paper, whilst limiting its analysis to age-eligible and non-age-eligible 3G households, nevertheless fails to demonstrate conclusively the differential impact of pension – rather than earned – income on labour supply of the marginal, non-earning working age individual. The reason is that it does not convincingly address the alternative causal factor: namely, processes of household formation.

Section 4, below, describes how we assess the different predictions of the two models using the PSLSD dataset.

## **4 Data and Empirical Methodology**

As noted in section 1, the focus of this paper is primarily on the manner in which BMM's results follow from their theoretical priors, the actual issue at stake – the true impact of the pension – is an important but secondary concern. As they and others have noted, the question of pension impact on household members is one that is inherently best answered with data that incorporates a substantial time component. Nevertheless, for the purposes of this paper – that is, to assess the validity of the BMM conclusions relative to alternative models applied to the same data - it is sufficient to closely follow the approach of BMM as Posel et al (2004) did.

### **DATA**

The data used is the 1993 Project for Statistics on Living Standards and Development (PSLSD) dataset.<sup>27</sup> It comprised 9,000 households that were surveyed in the latter half of 1993. Both for the reasons they outline – though note some of the concerns raised in Section 2 – and particularly for the purpose of replicating their results, we follow the approach of BMM. In particular, we focus on black African, three-generation households. Furthermore, for a large part of the analysis we will focus primarily on males of prime age (from the ages of 16 to 50) in those households, as this is the group for which BMM find the strongest results.

A number of individual- and household-level characteristics are used as control variables in the regressions. For purposes of replication and comparability the analysis utilises the same control variables as BMM – where additional controls are used these are noted in the text. In terms of dependent variables, whereas BMM use both labour hours and

---

<sup>26</sup> We test this possibility in section 5.

<sup>27</sup> BMM refer to this as the 'Integrated Household Survey of South Africa'.

employment status as their dependent variables, here we assess only the latter for two reasons: 1. The results for both dependent variables in the BMM papers are very similar 2. We are interested in attempting to differentiate between the HF and LS models using individuals that are not employed, for which hours worked are not relevant.

As with the BMM paper, we will also adjust the standard errors from our regression for correlations within PSLSD household clusters.

One issue noted in section 3 is that the HF model suggests that including total (per capita) income as a control in regressions estimating the impact of pension income may be important. However, for reasons given in the same section this is unlikely to be an issue with the static data in use here, and hence this control is included in only one of the regressions below (and it is found to be insignificant).

In terms of regression specification, because the results in the BMM paper, and our replications, are very similar in terms of sign and significance of the coefficients for the four primary measures of the pension – pension income, pension receipt, pension eligibility and pension income instrumented by the number of pensioners in the household – the analysis that follows, whilst favouring the IV estimation, does not put great weight on which of these is used.<sup>28</sup>

## **METHODOLOGY**

The empirical methodology is as follows. The first step is to (broadly) replicate the key results of the BMM papers. Thus, we replicate the results of BMM in which different pension variables appear to have a negative impact on labour supply. We then replicate their estimation of the impact of the same pension variables on various measures of employment status.

Next, we replicate the finding of Posel et al (2004) regarding the consequences of including non-resident members in the household size control variable. This is achieved by rerunning the base BMM regression with different household size variables.

Having replicated these results we then seek to test the hypotheses and implications of the HF and LS models, as outlined in sections 2 and 3. To begin, we look for evidence suggestive of household formation processes being induced by the social pension:

- Does pension receipt have a positive impact on the likelihood of a household member having moved into that household in the previous five years?
- Is the number/proportion of wage earners higher in non-eligible households, as predicted by the argument that 3G households may only exist if they are able to achieve a certain income level?<sup>29</sup>

---

<sup>28</sup> An exception is the case of the regressions estimating the impact of the pension on the propensity to be a migrant into the household, because in that instance the IV estimation produces significantly different results from the other specifications.

<sup>29</sup> This is the ‘existence’ argument of section 3.

- Relatedly, are there are significant differences between the per capita incomes of eligible and ineligible households? (The income smoothing HF model suggests that there should not be.)
- Lastly, are there structural differences in the reasons given across eligible and ineligible households for not engaging in job search? *If* the pension is acting as a safety net for the pensioner’s extended family, we expect that sick or disabled individuals are more likely to be present in eligible households.

The final step then is to examine the two key theoretical implications of the LS model derived in section 2:

- Limiting the sample to households already receiving a pension, does pension income continue to have a negative impact on labour supply as seems to be implied by the LS model?
- Do the various pension variables have a significant impact on whether an individual engages in job search *conditional* on their being unemployed?

## 5 Results

### REPLICATION

The key replication results are captured in Table 1 and Table 2. The former confirms the BMM results when regressing a binary employment variable on four measures of pension receipt. The first column of Table 1 is for all prime age individuals in 3G households, whilst the remaining columns are only for prime age males. Table 2 replicates the BMM finding that the pension does not have a significant effect on the probability of being *unemployed*, or of being *discouraged*, but does have a negative, marginally significant effect on *labour force participation*. (Note that these conclusions are challenged in alternative estimations discussed below).

In both these tables, and for all the instrumental variable specifications that follow, the coefficients on the IVs in the first-stage regression - of pension income on the number of pension-eligible men and number of pension eligible women – are highly significant.<sup>30</sup>

Table 3 confirms the finding of Posel et al (2004: footnote 14) that using a definition of household size that includes non-resident members results in the pension income variable being insignificant. The additional step that this paper takes is, in column (2), to introduce the number of non-resident members as a separate variable. In that instance the result is less extreme, although the coefficient on the pension variable remains insignificant at the 10% level. However, in the same regression conducted for the subset of prime age *males* – see column (3) - it is found that the result remains significant at the 5% level. This coheres to some degree with Posel et al’s main finding that the pension has a positive

---

<sup>30</sup> This is to be expected since in each case we are using the number of pension-eligible men and women to instrument for pension income. Although it is hypothetically possible that there may be subsets of the dataset for which this is not true, we do not find this to be a problem with any of the subsets used below.

impact on the employment propensity of *non-resident* young prime age women (in rural 3G households).

## MIGRATION

One key aspect of the BMM analysis that we fail to replicate is that which involves the impact of the pension on migration into a household ('positive migration'). Recall that migration in this instance is determined by the answer to the question, "Did \_\_ move here during the past 5 years".<sup>31</sup> BMM (2000: Table 10) find an almost significant, *negative* impact of the pension on migration of African prime age males, using pension income and the pension dummy as primary regressors. Using the eligibility dummy and the IV specification they then find an insignificantly positive effect.

Our results – see Table 4 – are importantly different. We also find a negative impact using the first two variables, but it is wholly insignificant. More importantly, the eligibility variable suggests a positive effect but is not significant at the 5% level. However, the preferred IV specification – column (4) for men and column (5) for men and women – yields the result that the pension *does* have a significantly positive effect on the probability of being a migrant, as predicted by the HF hypothesis.<sup>32</sup> Note in addition the highly significant association between propensity to be a migrant and the number of children (aged 0-5year or 6-15year) in the household. This is consistent to a degree with the findings of Edmonds, Mammen and Miller (2005) in which migration is by young adults *accompanied* by their children.<sup>33</sup>

## HOUSEHOLD FORMATION PROCESSES

Table 5 contains results of regressions aimed at exploring the association between pension eligibility and migration. In particular, it examines whether the pension is associated with non-employed prime age men having valid reasons for not engaging in job search (and, to some extent by implication, not being employed). Table 8 provides a tabulation of these reasons for men over the age of 30.<sup>34</sup> The proportion of this group that are sick and mentally disabled is notably higher in pension eligible households, but given the sample size these numbers can only be suggestive.

We create a binary 'good reason' variable equal to 1 if non-employed individuals had a good reason for not looking for work and 0 if they didn't; where disability (mental or physical), illness, being a housewife or engaged in 'child rearing' and saying that there

---

<sup>31</sup> PSLSD Household Questionnaire, p.6.

<sup>32</sup> Strictly speaking, the result in column (4) is not quite significant at the 5% level ( $p=0.055$ ) but we deem it to be close enough – especially given the result in (5).

<sup>33</sup> It is not *entirely* consistent with those results because here we find that the effect is significant for men and women whereas they only find significant results for women.

<sup>34</sup> The age restriction is due to the large number of younger men reporting being in 'formal education' as the reason for their not searching.

are ‘no jobs’ available all constitute good reasons. A similar ‘other reason’ variable is created, which excludes the ‘no jobs’ reason.

Table 5 presents the results of regressions in which these as used as the dependent variables. The coefficients on the pension variables for the ‘good reason’ regressions in columns (1) and (3) are contrary to that expected, but they are not at all significant. However, the sign on the coefficient in regression (2) – which primarily estimates the impact of the pension on the likelihood of a prime age male having health-related reasons for not engaging in job search – is as expected. Though not significant at the 5% level it suggests that an additional thousand Rand of pension income increases the probability of having health-related reasons for not searching (and by implication not being employed) by 4.4%. The difference between this result and that of regression (1) suggests that, in the HF framework, discouragement is not an important cause of in-migration. In order to test the veracity of such a conclusion we conduct an alternative estimation that uses pension income per capita, controlling for household per capita income, and excludes from the analysis those men who report being in ‘formal education’ as their reason for not searching. The coefficient for the OLS regression (4) is sizeably positive, but remains insignificant, whereas that for the IV estimation (5) is wholly insignificant.

Recall also that the basic HF model outlined in section 3 predicts similar per capita incomes across pension eligible and non-eligible households. Relatedly, it predicts (see page 12) that pension ineligible 3G households will have a greater number/proportion of wage earners. The data, and statistical tests thereon, support these assertions. Excluding outliers we find that average per capita household income is R165 and R162 for eligible and ineligible households respectively, and these are found to not be significantly different. In addition, on average *employed* prime age individuals comprise 12.5% of pension ineligible households but only 9% of eligible households. Ranking households by these proportions and using a Kruskal-Wallis test, eligible and ineligible households are found to be significantly different along this dimension ( $p=0.001$ ).<sup>35</sup>

## **EXTENSIONS OF THE LS MODEL**

The two key tests of the LS model proposed in section 2 relate to the marginal effect of pension income on labour supply conditional on pension receipt, and the impact of pension variables on the propensity to engage in search.

Table 6 seeks to address the first implication. Column (1) estimates the OLS specification for the subset of males in pension-receiving households.<sup>36</sup> Although the coefficient is negative in sign it is not at all significant. In other words, pension income does not appear to provide a disincentive to engage in job search as implied by BMM’s LS model. One

---

<sup>35</sup> We use the non-parametric Kruskal-Wallis test because the distribution of the variable in question is highly non-normal. See Keller and Warrack (2000: 602-607).

<sup>36</sup> It would make no sense to estimate the IV specification since the purpose of that was to deal with possible endogeneity in take-up, and in this instance we restrict our sample to those households that *are receiving* the pension.

concern might be that this result is due to the nature of the distribution of pensioners within households. Figure 2 confirms that this is heavily skewed to the right, showing that most of the variation therefore arises in the difference between households with one or two pension-eligible individuals. Some other work – for instance Duflo (2003) - has suggested that male pension income generally has less of an effect on other household members, and so we might be concerned that the non-effect of additional pension income is due to the limited distribution within the household of male pension income. In order to test this possibility, column (2) of Table 6 reruns the primary BMM regression for all prime age individuals, including male and female pension income separately. The coefficients on *both* are equally significant, suggesting that in our case this cannot explain the result in column (1).

The second implication of the LS model is tested in Table 7. We argued in section 3 that the one key area in which the predictions of the HF and LS model differ is in the propensity to search of unemployed men in pension eligible and ineligible households. Aside from being an additional test of the BMM hypothesis, recall that this potentially enables us to address one particular concern with the BMM estimations - replicated in columns (1) and (2) of Table 2; that is, these regressions include individuals *who are employed*, which from a HF model perspective not only adds unnecessary noise but may also distort the results.

The logic behind the estimation in Table 7 is thus as follows: Restrict the group of interest to only *those of working age that are not employed*; rerun the LPM regression of BMM with a binary dependent variable equal to 1 if searching unemployed and 0 if not in the labour force. The expectation, based on the HF model, is that pension eligibility will have a negligible impact on the dependent variable - assuming homogeneity amongst the unemployed in the two types of households. Whereas the LS model implies that the pension should have a negative impact on the dependent variable.

The results in Table 7 fail to support the LS hypothesis, with none of the estimations – some of which also control for household income – finding a significant negative impact of the pension on the propensity to search given that an individual is unemployed. Regression (2) is the same as the base regression, (1), but confined to those prime age unemployed males who do not report formal education as their reason for not engaging in search (and by extension, being unemployed).<sup>37</sup> As in Table 5 we also estimate a regression using per capita pension income and controlling for household monthly per capita income. Neither variable achieves significance at the 5% level, although the pension variable in (3) is significant at the 10% level ( $p=0.086$ ). As noted previously, our

---

<sup>37</sup> If an individual states that they are not engaged in job search because they are pursuing formal education it appears reasonable to assume that this is the reason they are not employed. Some caution is in order however since the causality could run in the opposite direction (i.e. unemployment leading to the choice to pursue formal education).

preferred specification is that in column (4) which instruments for pension income. In that estimation the coefficient on the pension variable is wholly insignificant.<sup>38</sup>

## 6 Discussion

This paper, contrary to the emphasis of Posel et al (2004) begins with the acceptance of the result that prime age men in pension eligible households have lower average labour supply than those in non-eligible households. The emphasis therefore is on assessing the veracity of different explanations of this fact. In particular, contrasting the labour supply (LS) model put forward by BMM, with the alternative household formation (HF) model outlined in section 3, using the PSLSD data.

It is important of course to emphasise that these models *need not* be mutually exclusive. As noted previously, it is possible that processes of household formation allow individuals to reduce their labour supply in response to pension receipt amongst their extended family – i.e. individuals may reduce their labour supply shortly before or after they migrate into a pension eligible household. In this sense at least we cannot rule-out the broader, negative, BMM hypothesis. To do so, and in order to capture the many nuances involved in both labour supply decisions over time, as well as household formation processes, one ideally requires recourse to panel data. Nevertheless, we do not seek to use such data – the Kwazulu Natal Income Dynamics Study (KIDS) for instance – here, because we wish to determine the validity of the assertions made by BMM *given* the cross-sectional, PSLSD data they use. In addition, it remains an open question whether panel data will enable us to make more confident assertions. For instance, panel data *per se* cannot differentiate between in-migration of an individual who was already unemployed versus one who left their job shortly before migrating, and therefore cannot distinguish between the more complicated relation between the HF and LS models referred to above.<sup>39</sup>

Leaving these broader concerns aside, *in the limited context of the PSLSD data* (to which BMM confine themselves) it *may* be possible to rule-out the BMM assertion that pension receipt is associated with labour supply reduction by those *already* in those households: By demonstrating a significant impact of the pension on migration into households; finding evidence of income smoothing across households; and/or showing there to be differences in the characteristics of non-employed individuals in these households.

The results presented in section 4 do, we claim, make a strong case - within the limitations - that migration based on income-smoothing was taking place amongst the households in the PSLSD survey. In the preferred IV specification of pension impact on

---

<sup>38</sup> Whilst we include the number of non-resident household members in these regressions, the results are not sensitive to this aspect of the specification.

<sup>39</sup> One would need reliable information on an individual's employment status prior to entering the household. This would ideally be based on a panel structure in which individuals are tracked when they migrate. More realistically it would entail questions on past labour force status; and arguably this would be as useful for a cross-sectional analysis as it would be for one based upon a panel.

propensity of being a migrant (Table 4, regressions (4) and (5)) we find a significant positive impact. In addition we find household per capita incomes to be statistically indistinguishable, and pension ineligible households to have a significantly greater proportion of employed prime age individuals.<sup>40</sup> Non-employed prime age men within eligible households are also more likely to have health-related reasons for not engaging in job search (Table 5 and 8) – which is suggestive of some kind of formation (‘negative’, if unemployed individuals remain in the household when sick, or ‘positive’ if sick family members migrate into eligible households) contributing to the lower average labour supply of this group.

An obvious extension of the LS model is that a marginal increase in pension income for a household *already receiving it* should have a negative impact on labour supply. Contrary to this assertion, the results in Table 6 find no such impact and as noted in section 5 this cannot be attributed to the distribution of pension income across households. Note that this finding may also be contrary to the positive formation hypothesis, in as much as we might expect higher levels of pension income to draw more family members needing support.<sup>41</sup> It remains consistent however with a *negative formation* process in which there are a limited number of possible dependents who can *remain* in the household for support.<sup>42</sup>

Whilst the evidence above counters BMM’s claim that migration is not a factor, as noted above it does not rule-out migration *associated* with a reduction in labour supply. To address that issue we examined the propensity to engage in job search conditional on being non-employed. Recall that the rationale is that if pension eligibility negatively affects labour supply in the manner hypothesised by BMM it should also, *ceteris paribus*, negatively affect the propensity to search given that an individual is without employment (see page 5). The results of Table 6 suggest that the pension has a negligible impact on search, which is consistent with the HF model but not with the LS model.<sup>43</sup>

In fact other work suggests that we might expect the pension to have a *positive* impact on search. Posel et al (2004) for instance suggest that pension income may support the job search of non-resident female household members. Our non-result may reflect the presence in pension receiving households of those who are unable to engage even in

---

<sup>40</sup> Notice that this is different from the result that we began with, namely that labour supply of prime age individuals is higher in non-eligible households. That is because here we are concerned with the supposition that either due to income smoothing, or conditions for the existence of 3G households, non-eligible households inherently require a greater proportion of wage earners.

<sup>41</sup> Although a simple norm such that one household only supports one extended family member and their dependent(s) would be sufficient to explain this result in such a context.

<sup>42</sup> So if the lower labour supply in eligible households is due to a sick child who because of the income chooses to stay in the household (when they otherwise would have left), additional income will not have any effect.

<sup>43</sup> One very important caveat to this, is that the PSLSD survey did not capture details on the *nature* of search. Thus we cannot address issues such as the fact that unemployed individuals in households with workers may use network search - and possibly therefore say that they have not searched for employment in the previous two weeks when for economic purposes we might think otherwise.

search (because they are ill or disabled), whose effect on the average propensity to engage in search is off-set by those individuals actually using pension income to search for employment. This reasoning is necessarily speculative, but note that neither hypothesis is consistent with the BMM model.

## 7 CONCLUSION

The evidence and reasoning presented above is highly suggestive of household formation processes being responsible for the lower average labour supply of prime age African males in three-generation, pension-receiving households. Given that we contrast the labour supply reduction model of Bertrand, Mullanaithan and Miller (2000 and 2003) with an alternative household formation model (see section 3) using the same PSLSD data, the definite assertion by these authors that policy makers should be concerned about evidence of negative consequences of the social pension appears at best premature.

Due to both data limitations, and the complex, intertemporal nature of the process under examination, the conclusions we are able to draw from the results in section 5, and discussion in section 6, are necessarily tentative. Nevertheless, contrary to the BMM model they do suggest that household formation is a contributory factor to the labour supply differential between eligible and ineligible households. By limiting the analysis to non-employed prime age individuals we attempt to eliminate the bias and extraneous noise that may be caused by such household formation processes. The insignificance of pension income for the job search decision is, we suggest, again at odds with the prior supposition that individuals within pension-receiving households use pension income to support their increased leisure.

Nevertheless, even if it is accepted that the real cause of the labour supply differential is household formation processes that make pension income into a ‘social safety net’, policy makers may still consider this to be a ‘perverse’ outcome of a grant ostensibly targeted at the welfare of the elderly. Either way, we would argue that the social welfare and economy-wide implications of this explanation are importantly different from one in which the pension effectively decreases the incentive to actually take-up available employment.

From a methodological perspective, the difference between the emphasis of the conclusions in this paper and that of BMM, reflects the importance of adopting *ex ante* theoretical priors. Whether one should ever favour one plausible prior over another remains a much-debated methodological concern. Yet the analysis above indicates the impact such preferences may have not only on analysis of an academic nature, but also on the more influential policy implications that follow. Unless there are very good reasons – preferably both theoretical and empirical – to do otherwise, it would thus appear prudent to adopt an agnostic attitude in such cases. After all, the very weight of conclusions derived from applied economics (and indeed empirical analysis generally) draws directly on its ability to differentiate between equally plausible hypotheses.

In short, whether income from the social pension causally reduces labour supply – rather than merely being associated therewith – amongst other household members remains an open question; though the results of this paper do appear to favour the latter hypothesis. To definitively answer this question will likely require fairly sophisticated panel surveys that follow individuals as they migrate, combined with more extensive models of the factors influencing household formation. Considering that there is already persuasive evidence of the *positive* impacts of the social pension, one wonders whether the effort this will require is justifiable except for any additional value it might have as an academic exercise.

## References

Alderman, H., Chiappori, P-A., Haddad, L., Hoddinott, J. and Kanbur, R. (1995). Unitary vs. Collective Models of the Household: Time to Shift the Burden of Proof? *World Bank Research Observer*, Vol.10.

Bertrand, M., S. Mullanaithan, and Miller, D. (2000). Public policy and extended families: Evidence from South Africa. NBER Working Paper 7594.

Bertrand, M., S. Mullanaithan, and Miller, D. (2003). Public policy and extended families: Evidence from Pensions in South Africa. *World Bank Economic Review*, 17(1): 1-25.

Binmore, K. (1993). *Game Theory and the Social Contract. Volume I: Playing Fair*. MIT Press.

\_\_\_\_\_ (1998). *Game Theory and the Social Contract. Volume 2: Just Playing*. MIT Press.

Cherry, T.L. (2001). Mental accounting and other-regarding behaviour: Evidence from the lab. *Journal of Economic Psychology*, 22: 605-615.

Duflo, E. (2003). Grandmothers and Granddaughters: Old Age Pension and Intrahousehold allocation in South Africa. *World Bank Economic Review*, 17(1): 1-25.

Edmonds, Eric, K. Mammen, and D. Miller (2005). Rearranging the Family? Household Composition Responses to Large Pension Receipts. *Journal of Human Resources* 40(1):186-207.

Galbraith, J.K. (1987). *Economics in Perspective: A Critical History*. Boston: Houghton Mifflin Company.

GCIS, Social Cluster Press Briefing, 11 May 2006.

Keller, G. and Warrack, B. (2000). *Statistics for Management and Economics*, Duxbury, CA.

Leonard, T.C. (2000). 'The Very Idea of Applying Economics: The Modern Minimum-Wage Controversy and its Antecedents' in Backhouse, R.E. and Biddle, J. (eds.) *Toward a History of Applied Economics*. Durham: Duke University Press, pp.117-144.

Mortensen, D.T. and Pissarides, C.A. (1999). 'New Developments in Models of Search in the Labour Market' in Ashenfelter, O. and Card, D. (eds.) *Handbook of Labour Economics, Volume 3*, pp.???

Posel, D., Fairburn, J. and Lund, F. (2004). *Labour Migration and Households: A Reconsideration of the Effects of the Social Pension on Labour Supply in South Africa*. University of KwaZulu-Natal.

Ranchod, V. (2006a). "Household responses to adverse income shocks: Pensioner out-migration and mortality in South Africa", Preliminary draft, University of Michigan.

\_\_\_\_\_ (2006b). The effect of the South African old age pension on labour supply of the elderly. *South African Journal of Economics*, 74 (4): 725-744.

Rawls, J. (1971). *A Theory of Justice*. Cambridge: Belknap Press.