The Financial Lives of the Poor: Theory and Evidence from South Africa

Erlend Berg

Dissertation submitted for the degree of

Doctor of Philosophy in Economics at

The London School of Economics and Political Science

July 2009

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

Declaration

I declare that the work presented in this thesis is my own.

Erlend Berg

Abstract

This thesis studies three aspects of the financial lives of poor people, using theory and empirical analysis of data from South Africa. The first chapter studies the effect of mode of payment on household savings behaviour. Savings have been linked to consumption smoothing, investment and economic growth. A theoretical framework is set up to study four possible mechanisms through which electronic payment may have an effect on savings. Using difference-in-difference estimates around a policy change, it is found that households are more likely to have savings when their grant is paid directly into an account. These findings are mapped back to the theory in an attempt to identify the relevant mechanisms. The second chapter analyses funeral insurance as a distinct form of insurance. Funeral insurance is probably one of the oldest and historically most important forms of insurance, and is still widespread in parts of Africa today. Though funeral insurance is often provided by informal risk-sharing groups, this analysis abstracts from organisational form and asks under what circumstances funeral insurance is preferred to general life insurance. A model is set up in which there is an inter-generational conflict of interest over funeral expenditure, which funeral insurance may resolve. Predictions from the model are consistent with results from empirical analysis. The third chapter investigates an assumption which is nearly ubiquitous in development economics — that of constrained household liquidity. Direct empirical evidence of such constraints is surprisingly scarce. Many observations consistent with liquidity constraints are equally consistent with precautionary saving or a lack of forward planning. Using a South African panel data set and a source of anticipated income, the standard model with perfect capital markets is rejected. Finding that it fails, further analysis attempts to distinguish between liquidity constraints, precautionary saving and myopia as possible explanations.

Acknowledgements

First and foremost I would like to thank my advisors, Tim Besley and Maitreesh Ghatak. I am grateful for their support and guidance, and for the opportunities they provided me with. I would also like to thank Oriana Bandiera and Robin Burgess for the discussions we had and for their helpful suggestions. Tim and Robin lectured the development economics course I took as a masters student, thereby triggering my interest in research in the field. I am also grateful to Jørn Rattsø at the Norwegian University of Science and Technology (NTNU) for his advice and support.

I thank all my friends and fellow researchers at the LSE, and STICERD/EOPP in particular, for discussions and comments on my work, and for creating such a stimulating research environment. I have also benefited enormously from the thriving community of Goodenough College.

Thanks are also due to Jonathan Leape and Rashad Cassim for facilitating my visit to the University of the Witwatersrand, Johannesburg, in 2005.

Financial support from the LSE, the Economic and Social Research Council and the Research Council of Norway is gratefully acknowledged.

I owe a lot to my friends for their support over the past few years. Niraj and Wenche deserve special mention.

I am also grateful for my family's support and understanding. I have not been able to spend as much time with them as I would like.

I dedicate this thesis to my wife Shipra, who makes it all worthwhile.

Contents

List of Figures

List of Tables

Preface

An important function of financial markets is to allow households to shift consumption between periods and across stochastic states. Credit, savings and insurance are three ways of achieving this: Credit can bring consumption forward in time, savings can push it back, and insurance can shift it between different possible future states. However, as Besley (1995*b*) makes clear, the three serve overlapping purposes. If credit is available, it can be drawn upon on a rainy day and hence provide a substitute for insurance. Savings can serve a similar function, a notion explored by Deaton (1991). Nor are they always clearly distinguishable: Udry (1994) discusses a context in which state-contingent loan repayments blur the boundary between credit and insurance.

Households do not act in isolation. Pooling the risks of many households is the basic mechanism underlying mutual insurance. And what is saved by one household may be lent to another, often via banks or other financial intermediaries. The role of these intermediaries is to channel savings to the best investment projects. The importance of financial markets in development therefore goes beyond consumption smoothing: If the markets are dysfunctional, many worthwhile projects are unlikely to be financed, and this in turn will stifle economic growth.

For these reasons, it is not surprising that household finance is seen as a key to development. But our understanding of day-to-day financial decisions of households in poor countries is still sketchy. This thesis looks at aspects of the financial lives of poor households. The three chapters are laid out along similar lines: Each of them looks at behaviour relating to credit, savings or insurance. A theoretical framework is presented, from which predictions are extracted. The predictions are then subjected to empirical analysis using data on South African households.

Encouraging households to save is typically seen as desirable, but it is not clear how to go about it. The first chapter looks at whether the mode of payment of social grants matters for savings. The findings presented here suggest that, holding the amount constant, it makes a difference whether a payment is made in cash or deposited in a bank account. A model of household saving is laid out and used to identify four classes of mechanisms by which mode of payment may matter for savings: First, it may increase the return on saving or reduce the associated transaction cost. Second, there may be an income effect due to the time saved by not having to queue up for the cash pay-out at a specific time and place. Third, it may help households with time-inconsistent preferences to control their impulse expenditure. And fourth, it may increase the bargaining power of women in the household and thereby change the household's joint savings decision.

In 2003, social grant beneficiaries in three of South Africa's nine provinces were offered bank accounts into which their grants would be paid. It was a voluntary scheme, so many people already receiving grants continued to take them out in cash. There is anecdotal evidence that new mothers signing up for a child-support grant

were especially targeted for the new scheme, partly because they were deemed to be younger and more receptive to the use of new technology than the other large group of grant holders, the old age pensioners.

Using difference-in-difference estimates around the policy change, it is found that households receiving a child support grant are more likely to have savings of some form when they have the option of having the grant paid directly into an account. The story seems to be confirmed by the finding that the bulk of the increase is in bank account savings as opposed to in other forms, and there is also no evidence that the results are driven by a portfolio reallocation rather than by fresh savings. Mapping the findings back to the theory, there are indications that an increase in the return on savings may explain at least part of the effect. Behavioural effects may also matter, but there is little evidence of an income effect or an effect linked to the increased bargaining power of women. The estimated effect is largely driven by households headed by women. This may be due to actual differences in behaviour between male- and female-headed households. But it is also possible that the gender difference is due to a measurement error: If women have savings but do not tell their husbands about them, then the savings will only be reported if the female account-holder is herself the head of the household.

At least since classical Greece, people have formed risk-pooling groups with the specific aim of covering the funeral expenses of group members or their close relatives. These funeral societies are well-known to historians, and are documented to have existed at some point in most parts of the world. Even today, they remain widespread in many developing countries. In Sub-Saharan Africa they may represent the most common form of explicit insurance contract. Given their importance, funeral societies have received remarkably little attention from anthropologists, sociologists or economists. Dercon, De Weerdt, Bold, & Pankhurst (2006) is one of the few exceptions, documenting funeral associations in Ethiopia and Tanzania. They highlight a remarkably consistent feature of these societies: The terms of the contract are explicit and well-known to the members even if not always written down, yet they operate outside the formal economy.

Clearly, the large literature on informal insurance is relevant to the study of funeral societies. But the main focus of that literature is on testing for full insurance, and explaining the shortfall in terms of their reliance on self-enforcing contracts. The study of the specific insurance product offered by funeral societies has been largely neglected. The aim of the second chapter is to start filling the gap by discussing funeral insurance as a distinct type of insurance.

The starting point is the observation that the primary purpose of a funeral society is to arrange dignified funerals for covered individuals, whereas the payout from a modern life insurance is typically freely disposed of by the beneficiary. The strategy followed here is to abstract from both organisational form and enforcement constraints, and ask under what circumstances funeral insurance is preferred to general life insur-

ance. A model is set up in which there is an inter-generational conflict of interest over funeral expenditure, which funeral insurance can help resolve. Predictions are drawn from the model, though the available data prevents clean testing of the model as a whole. The main prediction tested is that holding household income fixed, greater wealth (assets) is associated with a shift from funeral insurance to life insurance. The prediction is consistent with results from analysing a large, nationally representative survey data set.

That households face credit constraints is close to being a standard assumption in the theory of economic development. But most of the existing empirical work attempting to document household liquidity constraints in a rigorous manner rely on data from rich countries. It is paradoxical that there is very little unequivocal evidence of liquidity constraints from the developing economy contexts with which the constraints are most commonly associated. Furthermore, many observations described as violations of the standard model with perfect capital markets are arguably as consistent with precautionary saving, or a lack of forward planning, as they are with liquidity-constrained but forward-looking households.

In the third chapter, the generous South African social pension is used as a source of anticipated income to study the resulting response in household expenditure in a panel of households. In the standard model with perfect capital markets, a version of which is laid out here, an anticipated increase in income should not be associated with a jump in consumption. Hence the standard model with perfect capital markets is rejected by the finding that expenditure breaks out of its household-specific trend when a household member reaches the pension-qualifying age, controlling for other demographics. Finding that the standard model fails, further analysis attempts to distinguish between liquidity constraints, precautionary saving and short-sightedness as possible explanations. Though data limitations reduce the power of some of the tests, the findings are on balance more supportive of liquidity constraints than of precautionary saving or a lack of forward planning as explanations of the expenditure jump.

1 Does payment into bank accounts encourage saving?

1.1 Introduction

Economists and policy-makers have long been interested in the determinants of household saving, for several reasons. At the macro-level, Schumpeter was an early proponent of the idea that financial intermediation is a critical determinant of economic development. Under this view, policies that increase household savings and bring it into the formal financial sector are clearly of interest. At the micro-level, savings may be even more important in poor countries than in rich ones, because it allows households to smooth consumption and to invest even when insurance and credit markets are dysfunctional. And nowhere do these issues seem more pressing than in Sub-Saharan Africa. Its economic under-performance over the last decades compared to other developing regions has been linked to a lower savings rate and a higher rate of capital flight.

This paper asks what happened to savings when, in 2003, South African social grant beneficiaries got the option of having their benefits paid into a bank account rather than in cash. A simple model is set up and used to shed light on how the bank account programme might have an effect on saving in both single-person (unitary) and two-person (collective) households. It is shown that the two-person household case reduces to the single-person case but with a discount factor which is a function of the individual discount factors of the two household members. Difference-in-difference estimates of the effect of the programme on savings variables are computed using space and time variation in the availability of the bank account programme.

The main finding of the paper is that female-headed households are much more likely to have savings of some form when they have the option of having their child support grant paid electronically. Most of the effect is, as expected, related to savings in bank accounts. The finding is corroborated by supplementary analyses indicating that subgroups who were plausibly treated more intensively displayed stronger effects. In most specifications there is no significant effect on male-headed households. But it is not possible to determine whether the difference in estimated effect across household types (female- versus male-headed) reflects a genuine behavioural difference or a measurement problem related to imperfect information-sharing within the household.

This raises the question of how mode of payment affects saving. Four possible mechanisms are studied in the theory section. First, it may have been caused by an increase in the return on savings. By directly depositing the grant into a bank account, the recipient avoids the costs associated with a trip to the bank. These could include travel and opportunity costs as well as 'embarrassment' in front of a bank teller. It may also increase the expected return on savings by reducing the risk of

appropriation. The increasing returns story is tested in the paper, and the evidence seems to support it.

Second, there could be a wealth effect. South African cash grant beneficiaries spend a considerable amount of time collecting their money. Furthermore, the system is quite rigid in that the beneficiaries have to be present at the paypoint on a specific day each month or risk losing that month's grant. With the electronic payment programme, both of these problems were alleviated, so the beneficiaries should be unambiguously better off. But although the predicted effect on saving of making an additional transfer to the household in each period is positive, in the present setup it does not make a non-saver start saving or vice-versa.

Third, there is a possible 'behavioural' effect. It may be easier to control spending when one's money is in an account. For instance, someone who plans to spend a certain amount on household items may withdraw only that amount. Spending more would then require a separate trip to the ATM, incurring a cost which may be enough to deter an impulse purchase. The findings are consistent with a behavioural effect.

Fourth, the programme may shift the balance of power within the household. Since it is the woman who receives the bank account, her relative bargaining power may increase. But the theoretical predictions of this mechanism are not in line with the empirical results.

Comparing the empirical findings with the theoretical predictions, a reasonable conclusion is that the programme effect is most likely to be driven by an increase in the returns on savings, and/or by a behavioural story in which the bank account acts as a commitment device for savers with time-inconsistent preferences. However, it should be added that it is not difficult to think of variations of the model presented here in which an income effect could drive the observed results.

The main contribution of this paper is to demonstrate that mode of payment matters for household saving behaviour. This should be of interest to policy-makers, for at least three reasons: One is that paying social grants electronically can increase the proportion of households that save. This in turn will have beneficial effects on people's capacity to smooth consumption and to invest, particularly in the context of dysfunctional credit and insurance markets. Another is that the increase comes primarily in the form of savings in bank accounts. Under the widely held view that savings mobilisation and financial intermediation are important precursors to development, the form in which savings are held is critical. Finally, there is an ongoing debate on the link between financial exclusion and poverty. It has been claimed that people who don't have bank accounts (the 'unbanked') are deprived of many opportunities, and that access to a bank account can function as a 'gateway' to the formal sector of the economy and thereby help these people escape poverty. However, the exact mechanisms are sometimes unclear, and robust empirical evidence is scant. By showing that electronic payment can increase the households' propensity to save, the paper establishes one way in which 'being banked' can make a real difference.

The findings come with caveats. One is that the savings variables analysed here are binary. That is, they provide information on *whether* a household has savings, and on the form in which these savings are held, but not on amounts. Therefore, the empirical results concern what makes a household more (or less) likely to have savings, but is silent on whether households who already have positive savings save more or less as a result of the programme. Still, it is natural to assume that a programme which makes non-savers start saving can also make those who already save, save more, and the theory presented here is supportive of this interpretation.

Another caveat is that although the bank accounts programme was available to all social grant beneficiaries in the relevant provinces, there is evidence to suggest that it was marketed more intensely to young mothers receiving the child support grant. This fact will be exploited in the empirical analysis, but also makes it more difficult to extrapolate the findings to other demographic groups.

The rest of the paper proceeds as follows. The next section briefly reviews the related literature. Then some institutional background is provided on the child support grant and grant distribution system in South Africa. Thereafter a simple theoretical framework is set up and used to analyse the possible programme mechanisms, before describing the data and presenting the empirical strategy and results. Finally there is a brief discussion section.

1.2 Related literature

Gersovitz (1988) provides an overview of the literature on saving and development with a focus on the determinants of savings, both at the micro- and macro-level. In a later survey, Besley (1995*b*) reviews the literature on the relation between financial intermediation and growth. King & Levine (1993) find that the ratio of liquid liabilities of the financial system to GDP in 1960 is a robust predictor of GDP growth in the following three decades in a panel of countries. Attanasio, Picci & Scorcu (2000) is a more recent study that finds robust links between savings, investment and growth in a panel of countries.

Aryeetey & Udry (2000) provide an overview of saving in Africa, and point out the possible link between a low savings rate and economic under-performance. In their attempt at explaining why Africa has performed worse than other developing regions, Collier & Gunning (1999) emphasise Africa's poorly functioning credit markets and especially a severe problem of capital flight and a suboptimal flow of financing between the informal and formal sectors. It follows that an increase in formal savings could be beneficial for the economy as a whole whether it stems from an increase in overall savings or a portfolio re-allocation. Household savings may also be less prone to capital flight. An overview of recent research on household finance in South Africa is provided by Ardington, Lam, Leibbrandt & Levinsohn (2004).

There is a large literature on incentives to save and their responses in rich countries,

particularly on tax-advantaged retirement saving plans in the US. Poterba, Venti & Wise (1996) and Engen, Gale, Scholz, Bernheim & Slemrod (1994) summarise the empirical findings. Although there is no clear consensus on whether these schemes affect the amount of savings, it seems clear that they do have an effect on portfolio composition. Besley & Meghir (1998) predict limited usefulness of tax-based savings incentives in developing countries, raising the stakes for finding a savings-stimulating policy in that context.

The literature providing micro-level evidence on savings devices and incentives in developing countries is small but growing. Ashraf, Karlan & Yin (2006*a*) offer a random subsample of bank clients in the Philippines a savings account with a built-in commitment device. The funds are either inaccessible before a certain date, or they cannot be withdrawn until the balance has reached a certain amount. The authors find that people who display hyperbolic preferences in hypothetical survey questions are more likely to take up the product, and that the group which is offered the product saves significantly more than the control group. Anderson & Baland (2002) argue that rotating savings and credit associations (ROSCAs) provide a mechanism for women in Nairobi slums to safeguard savings from their husbands. Gugerty (2007) also looks at Kenyan data, and argues that ROSCAs may encourage saving by providing a commitment device.

There is also growing body of research on 'financial exclusion'. Claessens (2006) is an international overview, whereas Porteous & Hazelhurst (2004) focuses on the South African case. The importance of access to financial services is stressed, in part motivated by the perceived potential of bank accounts to serve as 'gateways' to other financial services and to the formal sector more generally. Bertrand, Mullainathan & Shafir (2004) discuss behavioural-economics insights in the context of poverty, and argue that small institutional shifts can lead to large changes in behaviour. They propose that making government transfers electronic by default may encourage poor people to open bank accounts, and suggest that "policies that encourage the take-up of bank accounts may naturally lead to improved savings and budgeting in general". Stegman (1999) also argues that a planned move to electronic payments of social benefits in the United States might increase savings amongst the recipients. Yet the body of robust empirical knowledge on whether and how access to financial services matters is very limited.

1.3 Background

Until recently, black South Africans were largely excluded from the formal financial system. During Apartheid, this was intentional policy on the part of the government. But even after the introduction of democracy it was hard to obtain fully-fledged bank accounts, because most banks required proof of formal-sector employment. As late as in 2003, only 48% of South Africans aged 18 and over had a personal bank account.

The proportion among black people would have been even lower, and many of these would be simple savings books as opposed to full transactional accounts, but even this represents a substantial increase compared to the situation in 1994 when democracy was introduced (Porteous & Hazelhurst 2004). The post office savings book was always available to everyone, but it was a very basic product and did not, amongst other limitations, allow direct crediting of transfers such as social grants.

1.3.1 The child support grant

The child support grant was introduced in 1998 to replace the state maintenance grant (Case, Hosegood & Lund 2005). The new grant was intended for primary caregivers (in practice mothers) of poor children. To qualify for the grant, the personal monthly income of the caregiver could not exceed ZAR 800 (≈USD 130) if the caregiver lived in formal housing in an urban area, or ZAR 1,100 otherwise. As of 2005, the means test limits had not been updated (Budlender, Rosa & Hall 2005).

The grant is payable per child, for an unlimited number of biological children and up to six non-biological children. In September 2002, the monthly grant value per child was ZAR 110. It was increased several times, and by September 2004 it had reached ZAR 170 (National Treasury 2005). When the grant was introduced, it was available for children up to and including the age of 6. From 1 April 2003 this was increased to 8 years, and from 1 April 2004 children up to and including 10 years of age were eligible (Rosa & Mpokotho 2004).

The grant is a substantial supplement to the income of poor households. As a benchmark, the legal minimum wage for full-time domestic workers for most of 2003 was ZAR 650 per month in rural areas and ZAR 800 in urban areas (Hertz 2005).

During the period studied in this paper, the grant was under rapid expansion. In April 2002, 1.9 million children were enrolled in the scheme. By April 2004, this number had increased to 4.6 million (National Treasury 2005).

The child support grant is one of a range of social grants which the South African government offers to individuals meeting certain criteria. At the end of the 1990s the old age pension was by far the most important of these, both in numbers of beneficiaries and amounts disbursed. Since then, the number of pensioners has essentially been stable, while the number of recipients of the disability and child support grants has increased rapidly.

1.3.2 The bank account programmes

Although eligibility criteria for the grants are the same across the country, their distribution was until recently the responsibility of the nine provincial welfare departments.[1] This gave rise to the institutional variation across provinces which this paper exploits.

[1] This changed in 2006 with the new centrally controlled South African Social Security Agency.

Before the programmes, a large majority of beneficiaries received their grants in cash.[2] Initially, grant disbursement was handled by government staff in most provinces, but by late 2002 all nine provinces had outsourced cash distribution to external contractors. In 2003, four contractors shared the nine provinces between themselves: PostBank, CPS, Empilweni and AllPay.

In August 2003, schemes were introduced whereby social grant beneficiaries in three of the provinces could open bank accounts and receive their grants by electronic transfer.[3] They differed somewhat in the details, but the main thrust of all of them was to pay the grants into accounts rather than in cash. The main expected benefit for the clients was one of convenience, as they would no longer be required to queue up at their designated cash payout point at a specific time every month in order to claim their money. Instead, the money would be paid into their account on the first day of the month, after which they could access it at any time and at any ATM. Arguably this was a better deal in Gauteng than in the other two provinces, since it is largely urban and has better ATM coverage.

From the banks' point of view, it was good business. The government paid a substantial fee for each grant they handled. In addition, the legal requirements of cash payment locations had become increasingly demanding in terms of seating space, toilet facilities, etc. By moving some beneficiaries over to electronic payments, the pressure on these facilities was eased. The cash payment 'caravans' had also been the targets of armed robberies on several occasions, so shifting cash management over to the existing ATM infrastructure will have reduced security costs substantially.

Although the scheme was open to beneficiaries of all government grants, women in receipt of the child support grant were the main targets of the programme in at least two of the three provinces. Not only were they the fastest growing group of grant beneficiaries, but compared to the pensioners they were younger and thus deemed to be more open to using technology such as debit cards, ATMs and PIN numbers. Also, compared to pensioners and disabled people, they were likely to be more pressed for time and therefore stood to gain more from time-saving and convenience associated

[2] There was always an option of having the grant paid into a bank account; the difficulty was obtaining an account in the first place. The banks typically required proof of formal-sector employment, which most beneficiaries did not have.

[3] PostBank offered accounts in the Northwest province, FNB and Standard Bank offered accounts in Eastern Cape, and AllPay offered 'Sekulula' accounts with ABSA Bank in Gauteng. In 2004, AllPay also started offering accounts in Eastern Cape.

In KwaZulu-Natal, and in Limpopo from 2004, CPS provided all beneficiaries with biometric smart-cards to aid identification and facilitate disbursement. By 2003 it was possible to store money on these cards, in other words withdraw less than the full amount of the grant in cash and leave the rest on the card. CPS had also developed a card terminal whereby vendors could accept payment directly from these smart-card. However, in an interview with the author in 2005 CPS said that the vast majority of beneficiaries withdrew the full amount of the grant in cash each month, possibly because they did not understand or trust the technology. The fact that very few vendors had invested in a card terminal also limited the usefulness of the technology. Also, beneficiaries were still required to be present at the paypoint each month for the grant to be loaded onto their cards and then withdrawn. Hence the CPS technology is quite far from a fully fledged bank account, and KwaZuluNatal and Limpopo are taken as 'control' provinces in this paper.

with the programmes.

It is important to consider why these particular provinces were chosen. Although it is not possible to provide a complete answer, in the Northwest and Eastern Cape provinces the programmes were initiated at least in part by the provincial welfare departments. AllPay chose to start its programme in Gauteng only, although they were distributing grants in four out of the nine provinces of South Africa at the time. This was presumably because it is predominantly an urban province and thus has a relatively good network of ATMs. It may also be because the pressure on the cash pay points was the greatest in this province. By September 2004, 26% of AllPay social grants in Gauteng were paid into bank accounts, and most of these were child support grants. In Eastern Cape, 8% of AllPay's grantees were paid via their bank accounts.[4] This type of information is not available for the other two contractors involved in the treatment.

1.4 Theoretical framework

To clarify thoughts, this section presents a simple theoretical framework. The aim is to shed light on how the programme may affect household saving through the following possible mechanisms. It may:

- increase the return on bank savings,
- reduce the time spent accessing the grants,
- help control impulse spending, and/or
- increase the bargaining power of the woman in a two-person household.

Assume there is a single consumption good. Households live for two periods and earn an income y in each period. There are two available savings products, A and B, with risk-free gross returns R_A and R_B, respectively. Borrowing is not possible. Let s_A and s_B denote the amount the household saves in each product, and define the total household savings variable $s = s_A + s_B$.

In what follows, the four possible mechanisms mentioned will be related to the optimal savings levels. But the data set used in the empirical section below provides only binary savings information. So the binary savings decisions (whether to save or not) are also of interest:

$$S = \begin{cases} 1 & \text{if } s > 0 \\ 0 & \text{if } s = 0 \end{cases}$$

$$\text{and} \quad S_i = \begin{cases} 1 & \text{if } s_i > 0 \\ 0 & \text{if } s_i = 0 \end{cases}, \quad i = A, B.$$

[4] Numbers obtained from AllPay.

Two types of households are considered: One-person households consist of a single male or female agent. Two-person households consist of one man and one woman. (There may be children, grand-parents and other household members present in both types of households, but it is assumed that they do not influence the level or allocation of savings.) Each agent maximises his or her lifetime utility from consumption,

$$u(c_1) + \delta_g u(c_2),$$

where c_t denotes the agent's private consumption in period t. The discount factor depends only on the agent's gender, $g \in \{m, f\}$. This gender difference will be discussed further below, but note that for simplicity the subscript will be dropped whenever the distinction is irrelevant.

All agents have identical CES instantaneous utility functions with elasticity $\alpha > 0$:

$$u(c) = \frac{1}{1 - \frac{1}{\alpha}} c^{1-\frac{1}{\alpha}}$$

1.4.1 The one-person household

The single agent's problem is to maximise lifetime utility from consumption subject to the intertemporal budget constraints

$$\begin{aligned} c_1 &\leq y - s_A - s_B \\ c_2 &\leq y + s_A R_A + s_B R_B \end{aligned}$$

and the non-borrowing constraints

$$s_i \geq 0, \quad i = A, B.$$

Since the preferences satisfy non-satiation, the budget constraints will bind. Substituting these into the expected utility, the problem becomes

$$\max_{s_A, s_B} u(y - s_A - s_B) + \delta u(y + s_A R_A + s_B R_B)$$

subject to

$$s_i \geq 0, \quad i = A, B.$$

Note again that both savings technologies are safe; uncertainty plays no role in the model.

The solution is as follows. (Calculation in the appendix.) Let $R = \max(R_A, R_B)$. Then the optimal amount of total savings is given by:

$$s^*(\delta, R) = \begin{cases} y \frac{1-(\delta R)^{-\alpha}}{1+R(\delta R)^{-\alpha}} & \text{if } \delta R \geq 1 \\ 0 & \text{otherwise.} \end{cases} \tag{1.1}$$

If the returns on the two products are equal, the agent is indifferent between them and the allocation of savings is indeterminate. Otherwise she will put all her savings in the product with the higher return.

Note that the expression for s^* is equal to the solution to the case with a single savings product with return R. In other words, the agent ignores the inferior savings product. This observation will simplify the following analysis.

The possible programme mechanisms will now be considered in turn.

The programme may increase the return on bank savings For most black South Africans, saving money in a formal bank account before the programme would have implied physically depositing cash in a bank branch or post office. Compared to this, the programme may have increased the return on account savings in two ways. First, by reducing the cost of saving. This includes the time and outlays associated with travelling to the bank to deposit cash, and also the possible psychological cost of dealing with bank tellers for people who may be illiterate or feel inferior (Bertrand, Mullainathan & Shafir 2004, Verhoef 2001). And though the technology to keep savings safe from burglars and robbers was always available through products such as post office savings books, the programme probably did reduce the cost of safekeeping.

Second, it may have made bank savings safer from 'appropriation' by making it more secret. Anecdotal evidence suggests that people sometimes hesitate to save because family members, friends and neighbours will ask for contributions if they know someone has spare money. With the programme, saving in a bank account is simply a matter of withdrawing less than the full grant amount, so it is virtually unobservable. A risk of appropriation would correspond to a lower return in the current framework without uncertainty.[5]

Next, consider what happens to savings when there is an increase in the rate of return. Starting with the total amount of savings and assuming for now that $\delta R > 1$ so that there are positive savings, one may write:

$$\begin{aligned}
\frac{\partial s^*}{\partial R} &= y\frac{\alpha(\delta R)^{-\alpha}R^{-1}\left(1+R(\delta R)^{-\alpha}\right)-\left(1-(\delta R)^{-\alpha}\right)(1-\alpha)(\delta R)^{-\alpha}}{\left(1+R(\delta R)^{-\alpha}\right)^2} \\
&= y\frac{(\delta R)^{-\alpha}}{\left(1+R(\delta R)^{-\alpha}\right)^2}\left\{\alpha\left(R^{-1}+(\delta R)^{-\alpha}\right)-(1-\alpha)\left(1-(\delta R)^{-\alpha}\right)\right\} \\
&= y\frac{(\delta R)^{-\alpha}}{\left(1+R(\delta R)^{-\alpha}\right)^2}\left\{\alpha-1+(\delta R)^{-\alpha}+\alpha R^{-1}\right\}
\end{aligned}$$

The sign is determined by the expression in the curly braces. When $\alpha > 1$ it is positive

[5] Though a reduced risk of appropriation is modelled here as an increase in the return on savings, the welfare consequences are not clear. First, if one takes a household welfare perspective, appropriations may represent a welfare-enhancing transfer even if it represents a private loss to the would-be saver. Second, even if the household as a whole is better off by increasing its stock of savings, it is not clear that a policy of reducing 'approprations' is welfare-enhancing. For example, it could interfere with an informal credit-based insurance scheme and thereby reduce global welfare.

for all R. For $\alpha < 1$ it is positive when R is small ($R \longrightarrow \frac{1}{\delta}$), but turns negative as R increases. Regardless of α, $\frac{\partial s^*}{\partial R}$ will approach zero as $R \longrightarrow \infty$. In fact from the expression for optimal savings above, one can see that its limits are

$$\lim_{R \longrightarrow \frac{1}{\delta}} s^* = 0$$

and

$$\lim_{R \longrightarrow \infty} s^* = \begin{cases} 0 & \text{when } \alpha < 1 \\ \frac{\delta}{1+\delta} y & \text{when } \alpha = 1 \\ y & \text{when } \alpha > 1 \end{cases} .$$

So savings are only guaranteed to be increasing with the rate of return when the elasticity of intertemporal substitution is larger than 1, otherwise there will first be an increase with R and then a decrease. At the boundary $\alpha = 1$ savings will converge to a fixed proportion of income.

Intuitively, this result can be understood in terms of a substitution effect and a wealth effect. The substitution effect results from the fact that when the interest rate increases, consumption today becomes more expensive relative to consumption tomorrow. But this will only tend to postpone consumption, and thereby boost savings, if the intertemporal substitution α is greater than one. Otherwise it will tend to reduce savings. The wealth effect on savings is unambiguously positive since the present value of lifetime income falls when the interest rate increases, reducing consumption today. The total effect is therefore positive when $\alpha > 1$, and ambiguous otherwise (Besley & Meghir 1998).

Besley and Meghir also review empirical estimates of the intertemporal elasticity of substitution. The weight of evidence seems to be that it is less than 1, especially in developing countries. In the framework presented this would imply that the effect on savings of an increase in the rate of return is ambiguous.

Returning to the question of allocation, let the return on product B (bank) increase from R_B to $\widetilde{R}_B$. The effect on the amount and allocation of savings will depend on the allocation of savings before the change. If the agent saved in A before the change, then either $\widetilde{R}_B < R_A$ in which case there is no change in amount or allocation, or $\widetilde{R}_B > R_A$ in which case savings are shifted into B and the amount changes according to the analysis for the single-product case above. However if the agent saved in B before the rate increase, then the savings stay in the bank but the amount will change as per the above.

If none of the products were attractive enough before the change ($\delta R < 1$), then either the new rate $\widetilde{R}_B$ is still not attractive enough and there are still no savings, or $\delta \widetilde{R}_B > 1$ and the agent now saves in B. (For simplicity only the cases where the rates are unequal are discussed; it should be clear that whenever the rates are equal there is an indeterminate allocation across the products).

Turning to the binary decision variables, it is clear from (1.1) that optimal savings

is positive whenever $\delta R > 1$. So while an increase in the bank return R_B can make a saver out of a non-saver, it will not make someone who already saves stop saving completely. Therefore, the effect on the optimal S^* is non-negative. Similarly the increased return may make people who do not save in bank accounts start doing so, but nobody who already saves in the bank will stop. So the effect on S_B^* is also non-negative, and the effect on S_A^* non-positive.

The programme may free up time for grant beneficiaries The programme may also reduce the amount of time spent accessing the grant. In many cases people queue a whole day to receive the cash payout. There are also costs associated with the inflexibility of the payouts: Unless the beneficiary is present at the paypoint on a fixed date, he or she will lose that month's grant. The programme is an improvement on both counts: When the grant is paid into the account, it can be accessed at any time from any ATM.[6]

Since this will benefit the agent equally in every period, in the present framework this effect is most readily conceptualised as an increase in income y. From (1.1) it is clear that whenever there are positive savings, they are proportional to y. Hence an increase in y will encourage further saving for agents who already save. But there will be no effect on portfolio allocation, and also no change for those who did not save before.

In terms of the binary decision variables it will have no effect: No non-savers become savers or vice versa, and there is no reallocation between the product types.

The programme may help control impulse spending It is also possible that the programme has 'behavioural' effects on saving. Though there are several possible such effects which might play a role (see Bertrand, Mullainathan & Shafir (2004) for a discussion), only one will be considered here: Control over impulse expenditures. Consider someone who knows she has a problem with time-inconsistent spending preferences. If all her money is in her purse, the barrier to making a spontaneous purchase is likely to be low. But if the money is in a bank account, she may be able to control her expenditure by only withdrawing just enough money for planned expenditures. An impulse purchase will then necessitate a separate trip to the ATM, and that may be costly enough to control the 'urge'.

To conceptualise this, a framework is needed in which the agent plans in advance how much to save. A 'zeroth' period, the planning period, is introduced before the two periods described above. In period 0, the agent plans her savings s_P, but actual

[6]Some beneficiaries live far from the nearest ATM and therefore need to spend a considerable amount of time and/or money travelling to withdraw cash. However, these beneficiaries are likely to live in remote areas without many facilities. It is therefore plausible that many of them would need to travel to a local centre regularly for other purposes (such as shopping or visiting a health service) in any case. If these centres have cash machines, and if such trips take place at least once a month, then the marginal cost of making a cash machine withdrawal is likely to be negligible.

saving will still happen in period 1. For simplicity, consider a single savings product which is attractive in the sense $\delta R > 1$; later it will be made clear that relaxing this is straight-forward. Time-inconsistent preferences are modelled in the usual way by introducing an extra discount factor $\beta < 1$ on future periods. Also assume that the agent has no income in period 0, and that $u(0) = 0$.

Without a commitment device, the agent's expected discounted future utility in period 0 is:

$$0 + \beta\delta u(y - s) + \beta\delta^2 u(y + sR)$$

The only difference between this and the expected utility in the non-behavioural case above is a multiplicative constant $\beta\delta$, so the optimal level of savings from that analysis is unchanged. Therefore, the agent sets $s_P = s^*$.

But in period 1, the myopic agent's expected discounted future utility is

$$u(y - s) + \beta\delta u(y + sR).$$

Compared to the non-behavioural case, this is a decrease in the discount factor. (Earlier it was δ, now it is $\beta\delta$.) Rewrite the expression for the optimal savings rate in the benchmark case as

$$s^* = y\frac{\delta^\alpha - R^{-\alpha}}{\delta^\alpha + R^{1-\alpha}}$$

and calculate

$$\begin{aligned}
\frac{\partial s^*}{\partial \delta} &= y\frac{\alpha\delta^{\alpha-1}\left(\delta^\alpha + R^{1-\alpha}\right) - \left(\delta^\alpha - R^{-\alpha}\right)\alpha\delta^{\alpha-1}}{\left(\delta^\alpha + R^{1-\alpha}\right)^2} \\
&= y\frac{\alpha\delta^{\alpha-1}R^{-\alpha}\left(R+1\right)}{\left(\delta^\alpha + R^{1-\alpha}\right)^2} \\
&> 0.
\end{aligned}$$

Intuitively, this is clear: The more highly you value consumption tomorrow relative to consumption today, the more you save.

Therefore, when the time comes to save (in period 1), the myopic agent will save less than she had planned (in period 0). If $\widetilde{s}$ denotes the optimal savings level from the point of view of period 1 in the absence of a commitment device, then clearly $\widetilde{s} < s_P = s^*$.

Now introduce a commitment device which makes it costly for the agent to deviate from her savings plan. In its simplest form, assume that there is a penalty F incurred in period 1 if the agent does not stick to her plan. It represents the cost of making an extra adjustment to the stock of savings in period 1, e.g. travelling to an ATM.

With the commitment device, the discounted utility in period 1 is

$$u(y - s - F \cdot 1\{s \neq s_P\}) + \beta\delta u(y + sR).$$

Here, $1\{s \neq s_P\}$ is an indicator function that is equal to 1 if the agent deviates from the savings plan, and 0 if she complies.

Clearly there is a critical level $\underline{F}(s_P)$ such that if $F > \underline{F}(s_P)$, the penalty outweighs the benefit to the agent of deviating from s_P. Assume first that $F > \underline{F}(s^*)$, so that s^* is enforceable. Then the agent will set $s_P = s^*$. To show this, note that the discounted utility in period 0 is now

$$0 + \beta\delta u(y - s - F \cdot 1\{s \neq s_P\}) + \beta\delta^2 u(y + sR).$$

Since the penalty is large enough to enforce s^*, the penalty will not be incurred when $s_P = s^*$. Therefore, the expected utility of setting $s_P = s^*$ has not changed compared to the case without commitment problems. Since $s_P = s^*$ is as attractive as before, and no other s_P can have been made more attractive by introducing a non-positive penalty, the agent will set $s_P = s^*$.

Since $s^* > \widetilde{s}$, the savings amount planned in period 0 is higher than the preferred (ignoring the penalty) savings amount in period 1. Therefore, the introduction of a sufficiently high penalty F will tend to increase savings (from $\widetilde{s}$ up to s^*).

Furthermore, since the penalty is linked to the bank account, an agent who has more than one savings product available, and who is sophisticated in the sense that she foresees her time inconsistency, will tend to shift her savings portfolio into the account with the penalty in order to safeguard the money from her future self.

Above it was assumed that $F > \underline{F}(s^*)$, so the penalty was sufficiently large to enforce s^*. But even if it is not, the device may allow a sophisticated agent to plan and enforce an intermediate amount such that $\widetilde{s} < s_P < s^*$. As long as $F > \underline{F}(\widetilde{s})$, there is an increase in savings associated with the device.

In terms of the binary variables, the commitment device may also turn non-savers into savers. By increasing the effective discount factor faced in period 1 from $\delta\beta$ to δ, myopic agents with $\delta\beta R < 1 < \delta R_B$ will shift from zero to positive savings, in the bank account, after the introduction of the device. Generally, the effects on S^* and S^*_B are non-negative while the effect on S^*_A is non-positive. But note that the device cannot turn a fundamentally unattractive savings technology, in the sense $\delta R < 1$, into an attractive one.

The above analysis applies to *sophisticated* time-inconsistent agents, in the sense of Strotz (1956). *Naïve* agents do not foresee that their time-inconsistency will continue in the future. Hence they will not see the impulse-control advantage of leaving their savings in the account as opposed to keeping it in cash, and the response in their savings behaviour will be weaker than for sophisticated agents. However, if they decide to keep their savings in the account for other reasons, such as higher returns in the account than for cash holdings, then they will benefit from the commitment device and save more than they would without the account.

1.4.2 The two-person (couple) household

In the empirical part of the paper it will be shown that the effect of the programme on saving is driven by female-headed households, so it is useful to think about why this might be the case. It may be a measurement issue: Spouses may lead separate economic lives with limited information about each other's savings. Since the grant is almost always paid to the child's mother, she will be the one whose saving is more directly affected by the programme. But the survey respondent is usually the household head. So in a male-headed household, the survey respondent may not know about any changes in the savings situation of his spouse. By contrast, in female-headed households the respondent will typically be the account holder herself. Therefore, the observed effect may be larger for female-headed households.[7]

An economically more interesting explanation of the gender heterogeneity in observed effects is related to intra-household decision-making. Men may have different preferences from women, and the household's saving may be thought of as the outcome of internal decision-making process. This section is devoted to studying the programme's effects on saving in a two-person household.

Household behaviour has often been modelled in essentially the same way as individual behaviour. Though the household might consist of more than one person, it was assumed that the intra-household decision processes would result in preferences that have the same properties as individual preferences. However, the literature has repeatedly rejected this framework, known as the unitary model, as a basis for explaining household behaviour.

One of the most promising alternative approaches to understanding household behaviour is the collective household model, see e.g. Vermeulen (2002). A major part of the appeal is its parsimony of assumptions. Essentially the only assumption needed is that the outcome of the household bargaining process is Pareto efficient. This assumption has for the most part not been rejected in the literature. In the case of South Africa in particular, Quisumbing & Maluccio (2003) test and reject the unitary household model but cannot reject that household allocations are Pareto efficient.[8]

The assumption that there is only private consumption will be retained. The opposite extreme, where consumption goods are public within the household, has been studied by Browning (2000). He also points out that an intermediate and more realistic assumption, where both private and public goods are allowed for within the household, is much more complicated because each household member would have more than one decision variable. From what is known of African household behav-

[7] In the data used, about 45% of respondent households are headed by a female. But in less than 3% of the cases does a female head of household report living with her spouse or partner. In contrast, about 60% of male household heads report living with their spouse or partner. So for practical purposes 'female-headed household' is synonymous with 'single mother' whereas male households head may live with or without a partner.

[8] However, Duflo & Udry (2003) reject Pareto efficient allocations within the household in their study from Côte d'Ivoire.

iour, arguably assuming purely private consumption assumption is more realistic than assuming purely public consumption.

A two-person household (couple) lives for two periods and has household income y in each period. What is not consumed in period 1 is saved for period 2 in savings products A and B as before. The household cannot borrow. The household members have identical instantaneous utility functions $u()$ of the CES type with parameter α, consume privately and maximise lifetime utility as before. Let c_{it} be i's consumption in period t.

Critically, assume that the only difference between the spouses is that the man is more impatient, so that

$$\delta_m < \delta_f.$$

Browning (2000) defends this assumption in the context of saving by pointing out that men are typically older and have shorter life expectations than their partners. The implication is that women tend to outlive their partners, and therefore put more weight on future consumption. There is also empirical evidence that men are more impatient, for instance Kirby & Marakovic (1996)[9].

Following the collective household literature, it is assumed that the outcome of the intra-household allocation process is Pareto efficient. The basic result of the collective model is that household's problem can then be written as a weighted average of the (net present) utilities of its members. In the current framework, the household solves

$$\max_{c_{f1},c_{f2},c_{m1},c_{m2},s} \mu\left[u(c_{f1}) + \delta_f u(c_{f2})\right] + (1-\mu)\left[u(c_{m1}) + \delta_m u(c_{m2})\right],$$

subject to the intertemporal budget constraints and non-negative savings:

$$\begin{aligned} c_{f1} + c_{m1} &\leq y - s_A - s_B \\ c_{f2} + c_{m2} &\leq y + s_A R_A + s_B R_B \\ s_A, s_B &\geq 0 \end{aligned}$$

The conventional interpretation of the utility weight μ is that it represents the relative bargaining power of the woman. For now, assume that μ does not depend on the other parameters of the model; later it will be argued that μ is likely to increase with the introduction of the programme.

Note that though there is only private consumption in this framework, it does allow for 'caring' preferences of a specific form. Introducing the net present utility of one person as an additive term in the net present utility of the other implies a household problem which reduces to the above expression, except that the bargaining

[9] This study looks at a sample of American college students. So although one can probably assume most of the men and women in the study are in a similar family situation (i.e., no dependents), there is no guarantee that the observed gender difference in patience remains the same in the presence of children. However, the age argument is still valid, and in the absence of evidence to the contrary it seems reasonable to assume that the gender difference persists.

power coefficient shifts in favour of the most 'cared for' member.

In the appendix it is shown that the solution to the household problem is

$$s^*\left(\widetilde{\delta}, R\right)$$

where 'the household discount factor' $\widetilde{\delta}$ is defined by

$$\widetilde{\delta}^{\alpha} = \frac{\mu^{\alpha}\delta_f^{\alpha} + (1-\mu)^{\alpha}\delta_m^{\alpha}}{\mu^{\alpha} + (1-\mu)^{\alpha}}$$

and $R = \max(R_A, R_B)$ as before. So essentially the solution is the same as for the single-agent household, but with a household discount factor which is a 'generalised weighted average' of the spouses' discount factors, with weights determined by the bargaining power parameter.

Therefore, the analysis from the case of the single-person household largely applies. The effect of an increase in R and y, and the introduction of a commitment device, will be exactly parallel to what was found above, only with a different discount factor.

The programme may increase female bargaining power However, there is a new mechanism to consider: The programme may increase the bargaining power of the woman, μ. It is the woman who has access to the savings device with higher return and/or less risk of appropriation, it is she who saves time in the payout queue, and it is she who directly gets access to what may function as a savings commitment device. All of these things may increase her bargaining power in the household.

It is clear from the expression for $\widetilde{\delta}$ above that it is increasing in δ_f and δ_m, and that its magnitude lies between that of the two spouses:

$$\delta_m < \widetilde{\delta} < \delta_f$$

Not surprisingly, $\frac{\partial\widetilde{\delta}}{\partial\mu}$ has the same sign as $\delta_f - \delta_m$ (shown in the appendix).

In other words, increasing the bargaining power of the woman will increase the household discount factor, bringing it closer to the individual discount factor of the woman. This in turn will tend to increase savings levels for saving households.

In the present model, an increase in μ has no effect on savings allocation. But stepping outside of the model, one would expect the extra savings to come in the form of bank savings since the increased bargaining power is directly linked to the availability of a private bank account.

Increasing the effective discount factor may turn non-saving households into savers if $\widetilde{\delta}R$ was less than unity before but becomes larger than unity with increased female bargaining power. The effect on S^* is therefore non-negative. It is also plausible that the effect on S_B^* is non-negative and that on S_A^* is non-positive, though this does not strictly follow from the model.

1.4.3 The difference in programme effect across household types

It is of interest to look for heterogeneous treatment effects across household types. The programme under study was linked to the child-support grant, which is almost always paid to women (mothers of young children). Ignoring children, grandparents and other relatives as economic decision-makers, there are basically two household types to consider: Single mothers and two-person (couple) households. Therefore, there are two potential sources of heterogeneous treatment effects: Two-person households have lower discount factors than single women, $\widetilde{\delta} < \delta_f$. Also, only two-person households are exposed to the potential effect of an increase in the woman's bargaining power μ.

The effective discount factor in the two-person household is intermediate between that of single women and single men. In particular, it is lower than that of households with a single female decision-maker. This will give a lower savings rate for the male-headed household (single men or couples) than for the female-headed household for any rate of return and income.

But whether the effect of the programme is smaller or larger for male-headed households is sensitive to the parameters. It is shown in the appendix that for positive household savings $s^*(\delta, R)$, that is when $\delta R > 1$, one may write

$$\frac{\partial^2 s^*}{\partial \delta \partial R} > 0 \quad \text{if} \quad \alpha < 1,$$

but for general α the sign is indeterminate. This means that male-headed households will have a lower savings rate for any rate of return R, but without making further assumptions it is not possible to say whether their response to an increase in R will be stronger or weaker than that of female-headed households.

Above it was noted that part of the increase in the return on savings might be due to a 'secrecy effect': By keeping the savings secret the risk of appropriation may be reduced. Although it has not been modelled explicitly, it should be noted that to the extent that husbands represent a significant appropriation risk, the increase in the return on savings associated with the programme might be larger for couples than for single females (but also less likely to be reported).

Since a larger proportion of income is saved in female-headed households, an increase in income y will lead to a larger increase in savings for them than for male-headed households.

The behavioural effect may also differ depending on the type of household, but again the direction is not clear. As shown in the appendix, $\frac{\partial^2 s^*}{\partial \delta^2}$ is of indeterminate sign, which means it is not possible to make a general prediction about what type of household will have a larger divergence between planned (in period 0) and preferred (in period 1) savings levels. Thus it is not clear whether the programme's possible behavioural effect would be larger for male- or female-headed households.

Since bargaining power is not a feature of single-person households, this effect would tend to show a larger increase in savings for two-person households.

1.4.4 Summary of predictions on binary savings variables

This section summarises the above findings with special emphasis on the binary savings variables. The aim is to identify predictions which can be tested empirically.

An increase in the bank return R_B may generally increase or decrease total savings s^*. But while it can make a saver out of a non-saver, it will not make someone who already saves stop saving completely. Therefore, the effect on S^* is non-negative. Similarly it may make people who do not save in bank accounts start doing so, but nobody who already saves in the bank will stop. So the effect on S_B^* is also non-negative, and the effect on S_A^* non-positive.

An increase in y may change the amount saved but will never make a saver out of a non-saver or vice versa, nor will it change anyone's allocation across the products. Therefore, there is no effect on either S^*, S_A^* or S_B^*. However, it is worth noting that this prediction is particularly sensitive to the simple framework used here. For instance, by introducing a subsistence requirement, that is a minimum consumption level below which the household would save nothing, it would be possible derive a set of conditions under which an increase in income would make the household start saving. Introducing a fixed transaction cost on savings could give similar results.

The introduction of a withdrawal penalty F can make a non-saver into a saver, and also make some people shift their savings into the bank account from elsewhere. But in the current framework it will not make anyone stop saving, nor move their savings out of the bank. Therefore, the effect on both S^* and S_B^* are non-negative, and the effect on S_A^* is non-positive.

The above applies to both single- and two-person households. Under the assumptions, female-headed households will generally have higher discount factors than male-headed households. This implies that they are more likely to be savers for any given interest rate. But depending on the prevailing rate of return R_B before the policy change, the reform could see both types becoming savers, only female-headed households becoming savers because the new rate is still not attractive enough for male-headed households, only male-headed households becoming savers if female-headed households already save, or none. The same applies for the commitment device, whereas an increase in y is not expected to have an effect on the binary variables for any household type.

The collective household model implies that the effective household discount factor should increase (become closer to the personal discount factor of the woman) if the programme increases female bargaining power. For single-person households, there is no such effect. As mentioned above, the model itself does not predict a change in savings allocation with bargaining power, but as the increased bargaining power is somehow linked to the bank account, one might expect the extra savings to come in the form of bank savings.

Table 1.1 summarises the predictions in terms of the expected programme effects

on the binary savings variables.

1.5 Data

The main source of data is the Labour Force Survey, a six-monthly series of surveys collected by Statistics South Africa. The September 2002 and September 2004 waves form a repeated cross-section data set.[10] The outcome variables are derived from the following question:

"Does this household, or a household member, own any of the following financial assets?

1 = Money in a savings account at a bank
2 = Savings in a stokvel [South African term for ROSCA]
3 = Savings in pension plan or retirement annuity
4 = Unit trust, stocks or shares
5 = Cash loans which are expected to be repaid
6 = Life insurance
7 = Other savings, specify"

For each of the seven options, the household could answer "Yes" or "No". There is no information on amounts. The primary interest is in two binary savings variables derived from this question: Whether the household has savings in a bank account, and whether the household has any savings at all. The household is taken to have savings in some form if the respondent has answered "Yes" to at least one of the options above. There is no data on savings amounts.

For simplicity, only African (black) households are considered, defined as households headed by a person of African ethnicity. There were 20,114 black households in the 2002 survey, and 21,725 in 2004. Descriptive statistics are shown in Table 1.2. It is promising that there are changes in the savings variables, but that the household demographics seem fairly stable across groups.

In some specifications only households in receipt of the child support grant will be considered. Descriptive statistics for this subsample are shown in Table 1.3. The first two rows, bank savings and any savings, foreshadow the regression results. For this subsample, there are also other significant difference-in-differences when the sample is compared across years and provinces, in age, gender, marital/cohabiting status and education of the household heads, indicating the possibility of non-random selection. The following analyses will attempt to overcome this problem.

The econometric identification below relies on computing difference-in-difference estimates between Gauteng and the control provinces before (2002) and after (2004) the bank accounts programme was introduced. It is therefore of interest to verify that Gauteng and the other provinces were on a 'common' trend in the period before the

[10] Though the LFS is intended as a rotating panel survey series, there was a complete break in September 2004; the panel rotation in effect 'started over'.

programme. Table 1.4 presents summary statistics and difference-in-difference estimates for 2000-2 for key variables. A common trend cannot be rejected for many of the key variables, although the difference-in-difference estimate is significantly different from zero for the age of the household head and one of the five education variables.

The provincial unemployment data for 2002 and 2004 are taken from Statistics South Africa (2005).

1.6 Identification and results

The basic approach is to estimate the difference-in-difference in savings associated with the programme. The change in probability of having savings between 2002 and 2004 for treated households will be compared to the equivalent quantity for the households that were not treated. Since it is not known which households took up the bank account programme, only the effect of being offered the programme ('intention to treat') can be identified.

The focus is on Gauteng as a treated province because it is known that the programme was widely marketed to households receiving the child support grant, and that the take-up was substantial by September 2004. This kind of information is not available for the Eastern Cape and Northwest provinces. In particular the worry is that the programmes there were not widely marketed. Therefore, these two provinces are excluded from the analysis. (In general the results below are robust to the inclusion of all three programme provinces for some specifications, but not for all.)

The remaining six provinces serve as controls: Western Cape, Northern Cape, Free State, KwaZulu-Natal, Mpumalanga and Limpopo.

1.6.1 Ordinary least squares

Begin by looking at the sub-sample of households receiving the child support grant. The difference-in-difference estimator is equivalent to the coefficient d in the following regression equation:

$$y_i = a + bA_i + cG_i + dA_iG_i + \beta X_i + \varepsilon_i$$

Here, A_i is a binary variable equal to 1 if observation i was made in 2004 (after the reform) and 0 if it was made in 2002 (before). G_i is equal to 1 if observation i is of a household living in the treatment province Gauteng, and 0 otherwise. Note that the variable is not indexed over t since panel data are not available. One advantage of computing the difference-in-difference through a regression is that it allows for a vector of control variables (X) to be included.

The results are shown in Table 1.5. Except for the column 6, only households who receive the child support grant are included in these regressions.

Column 1 is the basic regression without control variables. Households affected

by the reform are 11 %-points more likely to have savings of any form. By way of benchmark, this increase corresponds roughly to moving the household up by about ZAR 200 (≈USD 25) in monthly expenditure, or increasing the years of schooling of the household head by approximately 4.5 in the sample studied.

In column 2 control variables are included (see table caption for details). The estimate falls to about 10 %-points but remains significant.

Columns 3 and 4 indicate that these increased savings rates are indeed driven by bank account savings. In column 3, the bank savings dummy is the dependent variable. The specification is otherwise identical to the one in column 2, and as one might expect the magnitude of the effect increases. In column 4 the dependent variable is 'any savings except bank savings'. As expected the programme does not have a significant effect on this variable.

To check whether the estimated effect is really driven by the bank accounts programme, a 'placebo' regression is run in column 5. Instead of looking at the change between 2002 and 2004, this regression examines the change between 2000 and 2002. The aim is to rule out that the households receiving the child support grant in Gauteng were on a higher 'savings trend' compared to those in the other provinces. A positive effect in this regression would go against the finding that the programme has an effect on savings. The result is presented in column 5, and it seems clear that there was no such trend. The estimated coefficient of interest is negative and not significant at the 5% level. The smaller number of observations in this column is due to the large-scale increase in the take-up of the child support grant in the latter part of the period studied.

Even if child support grantees in Gauteng were not on a higher 'savings trend' compared to the control provinces, it could have jumped from 2002 to 2004 for reasons unrelated to the bank accounts programme. But unless there is a specific reason to believe that the jump applies only to households that receive the child support grant, one should observe it for the population at large. The regression in column 6 is similar to the one in column 2, except that all black households are included whether they receive the child support grant or not. There is no evidence of a general increase in savings in Gauteng.

Column 7 begins to look at whether the effect on male-headed households (in which the grant-receiving mother's spouse is present) is different from that on female-headed households (typically single mothers), by interacting the basic difference-in-difference variables with a dummy indicating whether the household is female-headed. Though the effect on male-headed households is still significant, it falls to 4 %-points. The effect on female-headed households is all of 13 %-points higher, and the difference is significant.

Thus the preliminary findings indicate that the programme does have an effect on the savings of households receiving the child support grant. It is not driven by an increasing trend for these households stretching back to 2000, nor is Gauteng generally

on a higher savings trend. There is indication that the increase in overall probability of having savings is driven by an increase in bank account savings, and the effect is stronger for female-headed than for male-headed households.

Note that in most of the regressions discussed above, the coefficient on 'Gauteng' is negative and significant. The literal interpretation of this is that in 2002, child-grant households in Gauteng were less likely than the their counterparts in other provinces to have savings. In the next section, however, it will be shown that the coefficient loses significance in the instrumental variable estimation. This may indicate that child-support grantees in Gauteng differ systematically from grantees in other provinces, providing further indication that the use of the instrumental variables technique is warranted.

1.6.2 Instrumental variables

Although suggestive, the above analysis suffers from a possible bias. Because the child support grant expanded so rapidly in the period studied (take-up tripled between September 2002 and September 2004), it is hard to rule out a priori that the set of households receiving the grant changed differentially across the provinces. Indeed the difference-in-differences computed in Table 1.3 are suggestive of this. This could violate the underlying assumptions of the OLS analysis, and bias the results. But one may attempt to get around this by using grant eligibility as an instrument for grant take-up.

The two main criteria for accessing the child support grant are being the 'primary caregiver' (in practice, mother) of a young child, and being sufficiently poor. To avoid potential endogeneity of material wellbeing, the focus will be on the former criterion. A household will be considered *age-eligible* for the child support grant in a given year if there is a child below the age limit for the child support grant in the household in that year. For 2002 the grant was paid to carers of children up to the age of 6, and by September 2004 carers of children up to the age of 10 years were eligible.

The standard assumptions required in order to use age-eligibility as an instrument for grant take-up are that age-eligibility predicts take-up, and that age-eligibility will not influence changes in savings decisions differentially across provinces, other than through take-up. The first will be verified by the data. The second assumption cannot be tested but seems plausible.

The starting point for the estimation is a triple-difference: Consider the increase in savings incidence from 2002 to 2004 for child support grantees in Gauteng. Subtract from this the corresponding increase for non-grant households in Gauteng. Finally, subtract from this difference-in-difference the equivalent quantity computed in the control provinces. The resulting estimate corresponds to the triple-interaction term in the specification below:

$$y_i = a + bA_i + cG_i + dC_i + eA_iG_i + fA_iC_i + gG_iC_i + hA_iG_iC_i + \beta X + \varepsilon_i$$

As before, A_i and G_i are binary variables indicating post-reform and Gauteng province observations. C_i is a new binary indicating that household i receives the child support grant. The triple-difference specification allows for variation in savings propensity across time and provinces, grantee vs. non-grantee households and all two-way interactions of these.

Now the endogeneity concern described above corresponds to a possible co-determination of the take-up of the child support grant and the error term. In attempt to deal with this, an age-eligibility binary variable is constructed to serve as an instrument for child-support grant take-up. In 2002, the age-eligibility variable is 1 for households whose youngest member is six years or younger, and 0 for the others. In 2004, it is 1 if a household's youngest member is 10 or younger, and 0 otherwise.

The results are presented in Table 1.6. Column 1 represents the uninstrumented regression for the complete sample of households. The coefficient on the triple-interacted term corresponds to a triple-difference estimate of the effect of the programme. The estimate of 5.7 %-points is significant. But here the worry about endogenous selection of grant beneficiaries remains.

Column 2 is a first-stage regression of grant take-up on age-eligibility. Based on this regression, a predicted value for child grant status is created. As expected, age eligibility is a powerful instrument for take-up.

The second stage is given in column 3. It corresponds to the base regression from column 1, but with the original grant take-up variable replaced by the predicted variable. Because the predicted take-up is itself an estimate, the standard errors are estimated by bootstrapping with 200 iterations. The effect of the programme becomes insignificant.

Columns 5–7 repeat the exercise for the subsample of female-headed households. The raw OLS estimate is 13.4% and significant. But this time the second-stage estimate is also significant, and the magnitude is almost unchanged at 13.1%. (Again, the standard errors were estimated by bootstrapping with 200 iterations.)

So it seems that the programme increases female-headed households' propensity to save by about 13 %-points, though it has no detectable effect on the sample as a whole. For the subsample of female-headed households, this increase in the likelihood of having any savings roughly corresponds to increasing the schooling of the household head by almost 9 years.

Standard two-stage least squares regression analyses were also run, where age eligibility and its interactions were used as instruments for child grant take-up and its interactions. The results were qualitative similar and significant for female-headed households.

One potential concern with the instrument used here is that if poor people have more children, then age-eligibility may affect savings other than through take-up of the child-support grant. In particular, the concern is that poor people may be more likely to be age-eligible because they have more children, and also save less, thus violating

the exclusion criterion. It is difficult to get around this by correlating age-eligibility with poverty measures based on income, consumption or assets, since poverty status may be partly determined by the number of children a household chooses to have. However, household demographics, including children aged 0-4 and 5-9, are controlled for in the above regressions. Also the notion that eligible households are poorer and hence less likely to save than other households is at least partly allayed by the fact that the correlation between having savings of any form and age eligibility in the sample is positive at .052.

More generally, Heckman (1997) has pointed out that the using instrumental variables in the context of programme evaluation is invalid whenever individuals select into the programme on the basis of unobservables in the outcome equation or variables correlated with these. Here, the relevant question to ask is whether unobservable factors which matter for the effect of the programme on savings may influence take-up of the child support grant after controlling for age-eligibility. For instance, assume that some households are more likely to start saving as a result of the bank-accounts programme than others because of an unobservable characteristic (say, 'financial astuteness'). If, within the subset of age-eligible households, financially astute households are more likely to take up the child-support grant than the others, then this defeats the instrumental variable strategy.

This is a subtle argument which makes it difficult to establish definitively the validity of an instrument. But note that any characteristic which is controlled for in the regression passes the test, so that if the influence of 'financial astuteness' on saving can be captured by education and education is controlled for, the strategy is valid. The instrument is also valid if financially astute people are generally more likely to save than other households, but have the same response to the bank accounts programme. The same applies if financially astute people are more likely to save as a consequence of the programme but this does not correlate with their decision about whether to apply for the child support grant.

Columns 4 and 8 show the results of reduced-form regressions, where actual take-up is ignored completely. Take-up is replaced by age-eligibility in all the interaction terms. Though this should alleviate any concern related to the instrumental variables estimation, a difficulty with this specification is that age-eligibility qualifies people for the child-support grant as well as for the bank accounts programme. So if the take-up rate changed differentially across provinces (as there is some evidence for in Table 1.2) then the resulting difference-in-difference estimate may confound the effect of the grant itself with that of the bank accounts programme. It is found that female-headed treated households were 3.7 %-points more likely to have savings of any form, but the estimate is only marginally significant.

In all these regressions, control variables for household demographics, the education of the household head and province dummies are included.

1.6.3 Stronger effect on households with babies

There are reasons to believe that the effect on households with very young children might be stronger than that on households with slightly older children. Anecdotes suggest that the bank accounts programme was 'pushed' when mothers first applied for child support grant after giving birth. As there is paperwork to do in any case when signing up for the grant, it seems an opportune time to try to convince people to take part in the new programme. Even from the beneficiary's point of view, the procedure may seem less burdensome when it is part of the general grant application process.

Children born after the introduction of the programme (August 2003) would have been 0 or 1 year old at the time of the September 2004 survey. Therefore, separate reduced-form regressions are run for two subsamples of households: Those where the youngest child was 0 or 1 year old, and those where the youngest child was 2 or 3 years old. It is expected that the programme effect is stronger for the former subsample.

The results are given in Table 1.7. There is no significant effect on either subsample as a whole in columns 1 and 2. But when coefficients are allowed to differ by gender of the household head in columns 3 and 4, a significant effect of 28.8% is found for the households with babies, and a much smaller and only marginally significant effect for the households with children in the age group 2–3. In columns 5 and 6 the following control variables are included: Age, education (binary variables indicating, respectively, any education, at least six years, at least nine years, at least 12 years, and higher education) and marital status of the household head, counts of the number of male and female household members in the age groups 0-4, 5-9, 10-19, etc for every decade of age up to 70+, dummies for household size (1, 2, 3, 4-6 and 7+ members) and province. Here it is found that there is a significant programme effect on the savings of female-headed households of both categories, but the estimated coefficient for the 'baby households' is about twice as large as the one for households with somewhat older children, and the difference is statistically significant.

These reduced-form regressions thus corroborate the main finding, by identifying the expected heterogeneous treatment effects.

1.6.4 Portfolio composition

It seems clear that the programme increased the likelihood of having savings of some form, at least for female-headed households. The initial analysis suggested that this increase stems from an increase in bank account savings. To verify this, an instrumental variables regression similar to those reported in Table 1.6 was run for each of the savings types recorded in the survey. Only female-headed households were included in this analysis.

The results are given in Table 1.8. Only the second-stage regressions are reported. Standard errors were estimated by bootstrapping with 200 repetitions. As expected,

the biggest sub-component in the savings increase appears to be bank savings. But there are also significant increases in the propensity to save in stokvels (ROSCAs) and life insurance.

None of the savings types are significant and negative, which speaks against a pure portfolio reallocation story. The results indicate that the programme led to an overall increase in savings, as opposed to a portfolio re-allocation, and that this increase was primarily in the form of bank savings. However, there is an important caveat to this conclusion: There is only binary information on savings in each of these forms. So the analysis would not be able to capture partial portfolio re-allocations, where amounts held in some of these asset classes were reduced without running them down completely.

1.6.5 A cost reduction story?

In the theory section it was shown that one possible mechanism by which the programme could increase savings was by increasing the return on savings, in particular by reducing the cost of depositing cash in a branch. If this is so, then households with high initial cost of saving should display a stronger effect from the programme. The aim of this section is to investigate that hypothesis.

The main formal savings product available to the general black population before the programme was the post office savings book. Therefore, distance to the nearest post office should provide a rough measure of the costs of saving before the programme. The surveyed households were asked: "How long, in minutes, does or would it take from here to reach the nearest of each of these facilities using the usual means of transport?

a. Food market
b. Public transport
c. Pre-Primary/Pre-school centre
d. Primary school
e. Secondary school
f. Clinic
g. Hospital
h. Post office or post office agent
i. Welfare office"

For each of the listed options, the possible answer options were 0–14 minutes, 15–29 minutes, 30–44 minutes, 45–59 minutes, 60 minutes or more, or "Don't know."

Consider the following stylised story: Before the programme, a grant recipient wishing to save regularly in a postbank account would have to 1) travel to the grant cash payout point once a month, and 2) travel to the post office once a month. After programme treatment, a client would no longer have to travel to receive the grant as it would be paid directly into her account. However, she would still have to travel to

an ATM at least once a month in order to withdraw some money.

There is no direct information on the distance to either cash payout point or ATM location, but they are both likely to be in the community centre. If it is assumed that the ATM and the cash payout point are equally distant from the household's dwelling, then the effect of the programme is to save one trip a month to the post office. Therefore, the programme effect should increase with distance to the post office.

To test this, reduced-form regressions (i.e., ignoring grant take-up) were run in which the reform effect was allowed to vary by distance to the nearest post office. 'Post-reform', 'treatment province' and their interaction, were each interacted with distance variables. The results are presented in Table 1.9.

One would expect the coefficient of interest to increase as the distance to the nearest post office increases. This is indeed what is found for the first three distance brackets. In column 1 the full sample is used. It is found that for households living less than 15 minutes away from a post office, the coefficient is negative and not significant. The effect for the next two distance brackets are sequentially higher, and significantly so. But for the fourth bracket there is significant drop which effectively brings the effect back down to zero. For the fifth bracket there is again a positive increase in the effect, bringing the effect up to a high level as expected. In column 2 only female-headed households are included. The pattern is similar, except that the last distance bracket only brings the reform effect back up to the level of the third distance bracket. Note that again it appears that the effect is stronger for female-headed households.

Overall the results support the transaction costs story, at least in part. Both regressions are in accordance with the story except for the fourth distance bracket, at which point the effect seems to 'dip' down to around zero before recovering to the expected level at the next bracket.

In addition to the increased return on savings mechanism, the behavioural mechanism may also have a heterogeneous effect with respect to distance. If the programme increases the propensity to save because beneficiaries only withdraw enough money for planned purchases, then the cost of additional withdrawals may increase with household distance from the nearest ATM. On the other hand, it is perhaps more plausible that the correct metric in this case is the distance between the ATM and the place where the temptation to buy on impulse is most likely to occur. If these are both located in a village or urban commercial centre, then there is less reason to expect the behavioural effect to increase with the household's distance to the nearest post office.

1.7 Discussion

Referring back to the predictions made in Table 1.1, it is clear that the general finding that the programme was associated with an increase in savings, primarily in bank accounts, is consistent with three of the proposed mechanisms. It could be explained by an increase in the return on bank savings, the use of the bank account as a com-

mitment device, and/or an increase in female intra-household bargaining power. On the other hand, the finding is not consistent with a pure income effect in this setup. So one may reject the hypothesis that the *only* effect of the programme is to make the grant beneficiaries better off in terms of saved queuing time and easier access to grants.

There is also some support for the claim that the effect increases with the expected preprogramme cost of saving, which strengthens the increased-return hypothesis.

It seems clear that the identified programme effect is driven primarily by female-headed households. However, in this setup the only mechanism that unambiguously predicts a differential effect across household types is increased female bargaining power, and its prediction is the opposite of what the empirics suggest. This weakens the bargaining-power mechanism as a potential explanation.

The increased-return and commitment-device stories predicted an indeterminate differential effect across household types. They are both potentially consistent with the gender effect. An increased return on savings could explain the observed findings if the increase is large enough to make previously non-saving female-headed households into savers, but not enough to attract male-headed households with lower effective discount factors. Similarly if the withdrawal penalty is enough to deter female-headed households from changing their savings plans but not enough to do the same for the more impatient male-headed households, this could also drive the observed effect.

But another possible explanation for the gender effect cannot be ruled out: There could be a measurement error due to an imperfect flow of information within the household. It is likely that in most cases the survey respondent is the household head. But it may be that only the grant recipient, the woman, knows about the savings kept in the bank account. Thus it is possible that increased savings will only be observed if the grant-receiving mother herself is the head of the household, which would explain why the effect can only be detected in these households. It is not possible to distinguish this mismeasurement bias from a true cross-gender difference in saving.

In summary, it seems likely that the effect is driven at least in part by an increased return on savings. The results are also consistent with a behavioural (commitment-device) effect. Participants probably also benefit from easier grant access, but in the current framework this has no predicted effect on savings and therefore cannot on its own explain the findings. However, as mentioned, it is not particularly difficult to think of plausible extensions of the model where an increase in income could have an effect on the binary savings variables. Though it may be present, any effect of increased female bargaining power will tend to work in the opposite direction from what was found here, so it is probably not the main mechanism. It may be that the discussed mechanisms drive the stronger effect for female-headed households, or it may be due to survey measurement error stemming from imperfect intra-household information flow.

There are some caveats to the findings. One is that it only looks at binary sav-

ings variables; clearly information on amounts would have provided for richer results. Hence the results on portfolio allocation should be treated with caution as it is not possible to identify a savings draw-down in a product unless it is emptied completely. Also, the focus of the analysis has been on households in receipt of the child support grant, and it is not possible to guarantee that other groups would show the same type of response.

There is also a concern related to what people perceive as savings. Imagine a household with very little money. If this was held in a bank account, they may have answered in the affirmative when asked whether they have savings. But it is important to ask whether the same amount, held in cash, would have been thought of as savings. If not, the variables used in this study would be mismeasured and the results potentially biased. But even if this were to be the case, the programme does seem to have an allocative effect in that bank savings do increase. Even if no new savings were engendered, driving savings into the formal part of the economy through financial intermediaries is important because it addresses the wedge between savings and investment that has been pointed to in the literature on African development.

Gauteng is an urban province with relatively good infrastructure. Therefore, one should be cautious in concluding that the identified effect of the programme in Gauteng could be replicated in the rest of the country, or in rural areas more generally.

In future research, it may be interesting to look at the longer-term effect on the savings variables. For example, one might look at whether the results hold up two, three or five years. In a follow-up paper to their original article on commitment savings devices, Ashraf, Karlan & Yin (2006*b*) find that two and a half years after the savings accounts under study were originally opened, the treatment group no longer had significantly more savings than the control group. The authors suggest four possible explanations, of which one is that it may be necessary for the bank to constantly 'market' the product for it to be used continuously. An advantage of the programme studied here is that it does not need to be re-marketed or re-issued to keep working. Once set up, electronic payments recur until the child graduates from the grant scheme. This may mean that the increased propensity to save may be a more long-lasting effect than the effect identified by Ashraf et al.

This paper has argued that how income is received matters for household saving. In principle it was always possible for beneficiaries of the child support grant to save money, either in a postbank savings account, in a stokvel (ROSCA) or at home. But the results show a strong increase in the incidence of savings associated with payment of the grant directly into a bank account, indicating that these other options were too costly, too risky or too easily accessible for impulse spenders. Electronic payment is put forward as a practical way of increasing intermediated household savings in developing countries. This is one concrete example where 'being banked' matters.

Table 1.1: Predicted programme effects on binary savings variables

Possible mechanism	Effect on binary variables			Differential effect across household types
	Any savings	Bank savings	Savings except bank	
Increased return on bank savings	Non-negative	Non-negative	Non-positive	Indeterminate
Income effect	None	None	None	None
Commitment device	Non-negative	Non-negative	Non-positive	Indeterminate
Increased female bargaining power	Non-negative	Non-negative?	None?	Effect **only** on male-headed households

Table 1.2: Descriptive statistics for the whole sample

	Gauteng 2004	Control 2004	Gauteng 2002	Control 2002	DID (a-b)-(c-d)
Any savings	0.619	0.525	0.548	0.483	0.0297
	[0.486]	*[0.499]*	*[0.498]*	*[0.500]*	*[.0184]*
Bank savings	0.568	0.377	0.481	0.367	.0774***
	[0.495]	*[0.485]*	*[0.500]*	*[0.482]*	*[.0133]*
Any savings except bank	0.293	0.359	0.271	0.296	-.0409*
	[0.455]	*[0.480]*	*[0.445]*	*[0.457]*	*[.0201]*
ROSCA savings	0.098	0.116	0.052	0.102	0.032
	[0.298]	*[0.320]*	*[0.222]*	*[0.302]*	*[.0362]*
Pension fund savings	0.095	0.103	0.114	0.087	-.0344**
	[0.293]	*[0.303]*	*[0.318]*	*[0.282]*	*[.0131]*
Unit trust savings	0.017	0.008	0.026	0.011	-.00644*
	[0.129]	*[0.090]*	*[0.160]*	*[0.104]*	*[.00274]*
Loans to others	0.019	0.021	0.029	0.022	-.00843***
	[0.137]	*[0.142]*	*[0.169]*	*[0.148]*	*[.00212]*
Life insurance savings	0.190	0.216	0.182	0.186	-.0222*
	[0.393]	*[0.412]*	*[0.386]*	*[0.389]*	*[.0114]*
Other savings	0.004	0.034	0.022	0.014	-.0383*
	[0.066]	*[0.182]*	*[0.147]*	*[0.116]*	*[.0163]*
Child support grant	0.154	0.288	0.053	0.106	-.0808***
	[0.361]	*[0.453]*	*[0.224]*	*[0.308]*	*[.0187]*
Age of head	42.7	46.8	42.7	46.6	-0.187
	[13.5]	*[16.4]*	*[14.0]*	*[16.2]*	*[.472]*
Female head	0.311	0.481	0.301	0.440	-.031**
	[0.463]	*[0.500]*	*[0.459]*	*[0.496]*	*[.0119]*
Head married	0.537	0.463	0.541	0.502	.0343*
	[0.499]	*[0.499]*	*[0.498]*	*[0.500]*	*[.0156]*
Head any education	0.907	0.754	0.914	0.770	0.00875
	[0.290]	*[0.431]*	*[0.280]*	*[0.421]*	*[.00978]*
Head at least 6 years of schooling	0.785	0.562	0.785	0.571	0.00763
	[0.411]	*[0.496]*	*[0.411]*	*[0.495]*	*[.0163]*
Head at least 9 years of schooling	0.531	0.343	0.537	0.329	-0.0202
	[0.499]	*[0.475]*	*[0.499]*	*[0.470]*	*[.0102]*
Head at least 12 years of schooling	0.280	0.163	0.277	0.163	0.00337
	[0.449]	*[0.370]*	*[0.447]*	*[0.370]*	*[.0102]*
Head has higher education	0.052	0.024	0.059	0.028	-0.00365
	[0.221]	*[0.153]*	*[0.236]*	*[0.164]*	*[.00262]*
Household size	3.1	4.1	3.3	4.1	-0.155
	[2.2]	*[2.7]*	*[2.4]*	*[2.8]*	*[.0862]*
Observations	2708	13753	2817	11744	

Standard deviations for columns 1-4 in brackets. Standard errors for the DID column are clustered by province.

Table 1.3: Descriptive statistics for households receiving the child support grant

	Gauteng 2004	Control 2004	Gauteng 2002	Control 2002	DID (a-b)-(c-d)
Any savings	0.531	0.498	0.409	0.483	.107**
	[0.500]	*[0.500]*	*[0.493]*	*[0.500]*	*[.032]*
Bank savings	0.455	0.324	0.336	0.334	.129***
	[0.499]	*[0.468]*	*[0.474]*	*[0.472]*	*[.0332]*
Age of head	44.0	47.3	46.1	47.8	-1.57**
	[13.2]	*[15.5]*	*[14.7]*	*[15.4]*	*[.619]*
Female head	0.392	0.572	0.477	0.543	-.113***
	[0.489]	*[0.495]*	*[0.501]*	*[0.498]*	*[.00542]*
Head married	0.600	0.550	0.497	0.556	.11***
	[0.490]	*[0.498]*	*[0.502]*	*[0.497]*	*[.0131]*
Head any education	0.876	0.708	0.899	0.725	-0.0147
	[0.330]	*[0.455]*	*[0.302]*	*[0.447]*	*[.019]*
Head at least 6 years of schooling	0.715	0.485	0.725	0.506	0.00344
	[0.452]	*[0.500]*	*[0.448]*	*[0.500]*	*[.0372]*
Head at least 9 years of schooling	0.443	0.263	0.430	0.250	-0.00761
	[0.497]	*[0.440]*	*[0.497]*	*[0.433]*	*[.0177]*
Head at least 12 years of schooling	0.189	0.090	0.087	0.088	.0939***
	[0.392]	*[0.286]*	*[0.283]*	*[0.284]*	*[.00691]*
Head has higher education	0.010	0.003	0.000	0.003	0.0108
	[0.097]	*[0.050]*	*[0.000]*	*[0.057]*	*[.00687]*
Household size	5.2	5.9	5.9	6.5	-0.187
	[2.3]	*[2.7]*	*[3.1]*	*[3.0]*	*[.105]*
Observations	418	3963	149	1244	

Standard deviations for columns 1-4 in brackets. Standard errors for the DID column are clustered by province.

Table 1.4: Descriptive statistics for households receiving the child support grant - common trend check

	Gauteng 2002	Control 2002	Gauteng 2000	Control 2000	DID (a-b)-(c-d)
Any savings	0.409	0.483	0.279	0.322	-0.0312
	[0.493]	*[0.500]*	*[0.451]*	*[0.468]*	*0.025*
Bank savings	0.336	0.334	0.209	0.221	0.0138
	[0.474]	*[0.472]*	*[0.409]*	*[0.416]*	*0.025*
Age of head	46.1	47.8	50.5	48.2	-3.980***
	[14.7]	*[15.4]*	*[14.8]*	*[15.2]*	*0.69*
Female head	0.477	0.543	0.628	0.606	-0.0885
	[0.501]	*[0.498]*	*[0.486]*	*[0.489]*	*0.048*
Head married	0.497	0.556	0.279	0.435	0.0960*
	[0.502]	*[0.497]*	*[0.451]*	*[0.496]*	*0.042*
Head any education	0.940	0.735	0.860	0.709	0.0529
	[0.239]	*[0.442]*	*[0.349]*	*[0.455]*	*0.029*
Head at least 6 years of schooling	0.765	0.516	0.698	0.482	0.0338
	[0.425]	*[0.500]*	*[0.462]*	*[0.500]*	*0.032*
Head at least 9 years of schooling	0.470	0.260	0.326	0.264	0.148***
	[0.501]	*[0.439]*	*[0.471]*	*[0.441]*	*0.026*
Head at least 12 years of schooling	0.128	0.098	0.128	0.085	-0.0130
	[0.335]	*[0.298]*	*[0.336]*	*[0.280]*	*0.011*
Head has higher education	0.034	0.013	0.035	0.025	0.0109
	[0.181]	*[0.113]*	*[0.185]*	*[0.157]*	*0.011*
Household size	5.9	6.5	5.9	6.4	-0.00682
	[3.1]	*[3.0]*	*[3.0]*	*[3.2]*	*0.069*
Observations	149	1244	86	398	

Standard deviations for columns 1-4 in brackets. Standard errors for the DID column are clustered by province.

Table 1.5: OLS results

Dependent variable	1 Any savings	2 Any savings	3 Bank savings	4 Any savings except bank	5 Any savings	6 Any savings	7 Any savings
Post-reform	0.015	.113***	0.0484	.14**	.17***	.0613***	.116***
	[.032]	*[.0228]*	*[.0276]*	*[.0423]*	*[.0311]*	*[.0137]*	*[.0254]*
Gauteng	-.0737*	-.124***	-0.0237	-.232***	-.118***	-.0958**	*-.112****
	[.0327]	*[.0132]*	*[.0154]*	*[.0214]*	*[.0179]*	*[.0382]*	[.0229]
Post-reform * Gauteng	.107**	.0953**	.113***	0.0525	-.0583*	0.0247	*.0407***
	[.032]	*[.0297]*	*[.0288]*	*[.0323]*	*[.0242]*	*[.0185]*	*[.0147]*
Female head							-0.0412
							[.0624]
Female head * post-reform							-0.00546
							[.0365]
Female head * Gauteng							-0.0232
							[.0472]
Female head * post-reform * Gauteng							.13***
							[.0311]
Controls	No	Yes	Yes	Yes	Yes	Yes	Yes
Period studied	2002-4	2002-4	2002-4	2002-4	2000-2 (placebo)	2002-4	2002-4
Households	Child Support Grantees	Child Support Grantees	Child Support Grantees	Child Support Grantees	Child Support Grantees	All	Child Support Grantees
Observations	5774	5774	5774	5774	1877	31022	5774

Standard errors, shown in brackets, are robust and clustered at the province level. The control variables used are the age, gender, marital/cohabiting status and education indicators (any education, 6+, 9+, 12+ years of schooling and higher education) for the household head, household size (indicators for 1, 2, 3, 4-6 and 7+ members), household demographics (the number of household members by gender/age groups 0-4, 5-9, 10-19, etc for every decade of age up to 70+), the provincial unemployment rate at the time of the survey, and province fixed effects. The unemployment rate was not available in 2000 so it's not included in column 5. *** significant at 1%, ** significant at 5%, * significant at 10%.

Table 1.6: IV results

Dependent variable:	1 Any savings	2 Child support grant	3 Any savings	4 Any savings	5 Any savings	6 Child support grant	7 Any savings	8 Any savings
Post-reform	0.0534**	0.158***	0.0312	0.0439	0.0786***	0.186***	0.0655*	0.0745**
	(0.019)	*(0.019)*	*(0.033)*	*(0.024)*	*(0.017)*	*(0.022)*	*(0.036)*	*(0.022)*
Gauteng	-0.102***	0.0659***	0.0593	-0.102***	-0.184***	0.0608***	0.0263	-0.180***
	(0.0080)	*(0.012)*	*(0.048)*	*(0.011)*	*(0.0068)*	*(0.015)*	*(0.070)*	*(0.011)*
Post-reform * Gauteng	0.0303	-0.0760***	0.0427	0.0307	0.0338*	-0.0899***	0.0423	0.0380
	(0.019)	*(0.019)*	*(0.031)*	*(0.024)*	*(0.017)*	*(0.022)*	*(0.031)*	*(0.022)*
Child support grant	0.00776		0.0950		0.0363		0.114	
	(0.022)		*(0.063)*		*(0.038)*		*(0.12)*	
Child support grant * post-reform	-0.0207		-0.000162		-0.0489		-0.0606*	
	(0.017)		*(0.033)*		*(0.025)*		*(0.035)*	
Child support grant * Gauteng	-0.153***		-0.0834		-0.125**		-0.137**	
	(0.023)		*(0.052)*		*(0.035)*		*(0.069)*	
Child support grant * post-reform * Gauteng	0.0569**		-0.00196		0.134***		0.131**	
	(0.018)		*(0.031)*		*(0.026)*		*(0.066)*	
Age-eligible		0.209***		0.0207		0.222***		0.0282
		(0.022)		(0.012)		*(0.020)*		(0.027)
Age-eligible * post-reform				0.00415				-0.0159
				(0.016)				(0.019)
Age-eligible * Gauteng				-0.0399**				-0.0495**
				(0.016)				(0.015)
Age-eligible * post-reform * Gauteng				0.000721				0.0365*
				(0.017)				(0.018)
Regression type	OLS	1st stage	2nd stage	Reduced form	OLS	1st stage	2nd stage	Reduced form
Sample	All	All	All	All	Female-headed	Female-headed	Female-headed	Female-headed
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	31022	31022	31022	31022	13472	13472	13472	13472

Counts of the number of male and female household members in the age groups 0-4, 5-9, 10-19, etc for every decade of age up to 70+, education level of household head (any education, 6 years or more, 9 years or more, 12 years or more and higher education) and province dummies, are included as controls. Standard errors, shown in brackets, are clustered at the province level for columns 1, 2, 4 and 5. For columns 3 and 6 they are bootstrapped with 200 repetitions. *** significant at 1%, ** significant at 5%, * significant at 10%.

Table 1.7: Reduced-form regressions by the age of youngest household member

Dependent variable: Any savings	**1**	**2**	**3**	**4**	**5**	**6**
Post-reform	0.0347	0.0333	0.0303	.0411*	0.0307	.0689***
	[.0194]	*[.0193]*	*[.0275]*	*[.0184]*	*[.0263]*	*[.0155]*
Gauteng	.0524***	.109***	.0854***	.096***	-.0987***	-.0287*
	[.0139]	*[.0185]*	*[.0202]*	*[.0243]*	*[.0139]*	*[.0143]*
Post-reform * Gauteng	0.0318	0.00129	-.0574*	-0.0171	0.0123	-.0707***
	[.0194]	*[.0193]*	*[.0275]*	*[.0184]*	*[.0246]*	*[.0153]*
Female head			-.153***	-.1*	-.0795*	-0.0253
			[.0329]	*[.0461]*	*[.0355]*	*[.0349]*
Female head * post-reform			0.0181	-0.00511	0.0236	-0.00491
			[.0338]	*[.0189]*	*[.0376]*	*[.0222]*
Female head * Gauteng			-.196***	-0.00366	-.139***	-0.0664
			[.0329]	*[.0461]*	*[.0294]*	*[.0349]*
Female head * post-reform * Gauteng			.288***	.039*	.209***	.0975***
			[.0338]	*[.0189]*	*[.0371]*	*[.0172]*
Age of youngest household member	0-1	2-3	0-1	2-3	0-1	2-3
Control variables	No	No	No	No	Yes	Yes
Observations	4027	3536	4027	3536	4027	3536

The results indicate that the programme effect is stronger for households where the youngest child is 0 or 1 year old. This corroborates the story that mothers of babies born after the introduction of the programme were primary targets of the programme, possibly because they in any case needed to do some paperwork in order to apply for the grant. Standard errors, shown in brackets, are clustered at the province level. Control variables are the age, education (binary variables indicating, respectively, any education, at least six years, at least nine years, at least 12 years, and higher education) and marital status of the household head, counts of the number of male and female household members in the age groups 0-4, 5-9, 10-19, etc for every decade of age up to 70+, dummies for household size (1, 2, 3, 4-6 and 7+ members) and province. *** significant at 1%, ** significant at 5%, * significant at 10%.

Table 1.8: Portfolio composition

Dependent variable	1 Bank savings	2 ROSCA savings	3 Pension savings	4 Unit trust	5 Loans outstanding	6 Life insurance	7 Other savings
Gauteng	.0764**	-0.0224	0.00999	.00578*	-0.00447	0.0322	0.00526
	[.0361]	*[.0255]*	*[.0129]*	*[.00337]*	*[.00387]*	*[.0932]*	*[.0153]*
Post-reform	.0633***	0.0168	0.0169	-0.00322	0.00177	0.0245	0.0266
	[.0144]	*[.0384]*	*[.0141]*	*[.00518]*	*[.00256]*	*[.027]*	*[.0216]*
Post-reform * Gauteng	0.021	0.046	-.0573**	0.00182	-.00737*	-0.0264	-0.0356
	[.0177]	*[.0377]*	*[.0286]*	*[.00373]*	*[.00386]*	*[.0239]*	*[.0265]*
Child support grant	-0.0256	0.0902	-0.00895	0.00814	-0.0235	0.0503	0.0217
	[.059]	*[.0589]*	*[.0483]*	*[.0138]*	*[.0166]*	*[.0698]*	*[.032]*
Gauteng * child support grant	-.11*	-.0987**	0.00427	-.0132**	0.0152	0.0101	.0308*
	[.0619]	*[.0491]*	*[.0108]*	*[.0067]*	*[.0127]*	*[.0159]*	*[.0177]*
Post-reform * child support grant	-0.0793	-.0765***	-0.0158	-0.00616	0.00345	-.0383**	0.00226
	[.0496]	*[.024]*	*[.0218]*	*[.00438]*	*[.00967]*	*[.0192]*	*[.0194]*
Post-reform * Gauteng * child support grant	.211**	.087**	.0812*	-0.00754	-0.00942	.0802**	-0.0236
	[.107]	*[.0431]*	*[.0424]*	*[.00565]*	*[.00944]*	*[.0395]*	*[.0195]*
Observations	13472	13472	13472	13472	13472	13472	13472

Second stage IV regressions with take-up of the child support grant predicted by age eligibility. Standard errors, shown in brackets, are bootstrapped with 200 repetitions. The number of household members in the age groups 0-4, 5-9, 10-19, etc for every decade of age up to 70+ separately by gender, and province dummies, are included as controls. *** significant at 1%, ** significant at 5%, * significant at 10%.

Table 1.9: Exploring the transaction cost story

Distance to nearest post office/agent	1 Any savings	2 Any savings
All distances	-0.0296	-0.0326
	[.0221]	*[.0231]*
15 minutes or more	.0459**	.0951***
	[.0148]	*[.0111]*
30 minutes or more	.0209**	.0487***
	[.00843]	*[.00894]*
45 minutes or more	-.0783***	-.164***
	[.0157]	*[.0149]*
60 minutes or more	.144***	.168***
	[.0278]	*[.0279]*
Households	All	Female-headed only
Observations	30852	13397

Savings regressed on reform effect interacted with distance to nearest post office or post agent. If the effects of the reform are driven by transport costs alone, one would expect the estimated interaction terms to increase with distance. Only the coefficients on the triple-interaction terms (post-reform * treatment province * distance) are shown, but the double interactions (post-reform * distance and treatment-province * distance) were also included for each of the distance bands, along with the control variables: Age, education (any education, at least six years, at least nine years, at least 12 years, and higher education) and marital status of the household head, counts of the number of male and female household members in the age groups 0-4, 5-9, 10-19, etc for every decade of age up to 70+, and dummies for household size (1, 2, 3, 4-6 and 7+ members) and province. Robust standard errors, shown in brackets, are clustered at the province level. *** significant at 1%, ** significant at 5%, * significant at 10%.

A Proofs

A.1 Optimal savings in the single-agent household case

The Lagrangian function corresponding to the problem is

$$L(s_A, s_B, \gamma_A, \gamma_B) = u(y - s_A - s_B) + \delta u(y + s_A R_A + s_B R_B) + \gamma_A s_A + \gamma_B s_B,$$

where γ_i are the Lagrange multipliers corresponding to the inequality constraints. The solution satisfies

$$\begin{aligned} -u'(y - s_A - s_B) + \delta R_i u(y + s_A R_A + s_B R_B) + \gamma_i &= 0, \quad i = A, B \\ \gamma_i s_i &= 0, \quad i = A, B \\ \gamma_i &\geq 0, \quad i = A, B \\ s_i &\geq 0, \quad i = A, B \end{aligned}$$

First consider the case where positive amounts are held in both savings products. This implies $\gamma_A = \gamma_B = 0$ and so

$$\frac{u'(y - s_A - s_B)}{u'(y + s_A R_A + s_B R_B)} = \delta R_A = \delta R_B.$$

So both products are used only when the associated returns are equal. Total savings $s = s_A + s_B$ can be determined as if there were a single product with return $R = R_A = R_B$, but the allocation between the two products is indeterminate. (The client is indifferent as the products are identical.)

Next, if the client does not save at all one may write

$$\frac{u'(y - 0 - 0)}{u'(y + 0 + 0)} = 1 \leq \delta R_i, \quad i = A, B.$$

So when none of the two products provides a return which 'beats' the discount factor, nothing is saved.

Finally, consider the case where one non-negativity constraint is binding and the other is not. Assume that $s_A > 0$ and $s_B = 0$. This implies

$$\frac{u'(y - s_A - s_B)}{u'(y + s_A R_A + s_B R_B)} = \delta R_A \geq \delta R_B.$$

So when one product gives a better return than the other, the agent saves as if only the product with the higher return were available. She saves nothing in the product with the lower return.

To summarise, define $R = \max(R_A, R_B)$. The client will save nothing if $\delta R < 1$

and a positive amount if $\delta R > 1$. The total amount saved is given by:

$$\begin{aligned}
\frac{u'(y-s)}{u'(y+sR)} &= \delta R \\
\frac{(y-s)^{-\frac{1}{\alpha}}}{(y+sR)^{-\frac{1}{\alpha}}} &= \delta R \\
\frac{y+sR}{y-s} &= (\delta R)^{\alpha} \equiv \theta \\
y+sR &= y\theta - s\theta \\
s &= y\frac{\theta-1}{R+\theta} \\
&= y\frac{1-\theta^{-1}}{1+R\theta^{-1}} \\
&= y\frac{1-(\delta R)^{-\alpha}}{1+R(\delta R)^{-\alpha}}
\end{aligned}$$

where the temporary variable θ has been introduced.

If $R_A = R_B$ the agent is indifferent between the products and allocation is indeterminate, but otherwise she will put all her savings in the product with the higher return.

A.2 Optimal savings in the two-person household

The first thing to realise is that the household will not use a savings product with inferior return. So the household will behave as if there were a single product with return $R = \max(R_A, R_B)$. Non-satiation guarantees that the budget constraints are binding. Using them to eliminate c_{m1} and c_{m2} from the system, the problem becomes

$$\max_{c_{f1},c_{f2},s} \mu\left[u(c_{f1}) + \delta_f u(c_{f2})\right] + (1-\mu)\left[u(y-s-c_{f1}) + \delta_m u(y+sR-c_{f2})\right]$$

$$\text{s.t. } s \geq 0.$$

The corresponding Lagrangian function is

$$L(c_{f1}, c_{f2}, s, \gamma) = \mu\left[u(c_{f1}) + \delta_f u(c_{f2})\right] + (1-\mu)\left[u(y-s-c_{f1}) + \delta_m u(y+sR-c_{f2})\right] + \gamma s,$$

where γ is the Lagrange multiplier for the non-negative savings constraint. The solution must satisfy:

$$\begin{aligned}
\mu u'(c_{f1}) - (1-\mu)\, u'(y-s-c_{f1}) &= 0 \\
\mu \delta_f u'(c_{f2}) - (1-\mu)\, \delta_m u'(y+sR-c_{f2}) &= 0 \\
(1-\mu)\left[-u'(y-s-c_{f1}) + \delta_m R u'(y+sR-c_{f2})\right] + \gamma &= 0 \\
\gamma s &= 0 \\
\gamma &\geq 0 \\
s &\geq 0
\end{aligned}$$

When $s > 0$ this yields a system of three equations in three unknowns:

$$\begin{aligned}
\frac{u'(y-s-c_{f1})}{u'(c_{f1})} &= \frac{\mu}{1-\mu} \\
\frac{u'(y+sR-c_{f2})}{u'(c_{f2})} &= \frac{\delta_f}{\delta_m}\frac{\mu}{1-\mu} \\
\frac{u'(y-s-c_{f1})}{u'(y+sR-c_{f2})} &= \delta_m R
\end{aligned}$$

Rearranging:

$$\begin{aligned}
\frac{y-s-c_{f1}}{c_{f1}} &= \left(\frac{1-\mu}{\mu}\right)^{\alpha} = a \\
\frac{y+sR-c_{f2}}{c_{f2}} &= \left(\frac{1-\mu}{\mu}\frac{\delta_m}{\delta_f}\right)^{\alpha} = b \\
\frac{y-s-c_{f1}}{y+sR-c_{f2}} &= (\delta_m R)^{-\alpha} = d
\end{aligned}$$

where the temporary quantities a, b and d are introduced. Solve the first two equations in terms of c_{f1} and c_{f2}, and divide the first by the second to simplify the third equation:

$$\begin{aligned}
c_{f1} &= \frac{y-s}{1+a} \\
c_{f2} &= \frac{y+sR}{1+b} \\
\frac{ac_{f1}}{bc_{f2}} &= d
\end{aligned}$$

Substitute the solutions for c_{f1} and c_{f2} into the third and simplify to find

$$\frac{y-s}{y+sR} = \frac{1+a}{1+b}\frac{b}{a}d = ed,$$

where another temporary quantity e has been introduced. Now solve for s to find

$$\begin{aligned} s &= y\frac{1-ed}{1+eRd} \\ &= y\frac{1-e\left(\delta_m R\right)^{-\alpha}}{1+eR\left(\delta_m R\right)^{-\alpha}} \\ &= y\frac{1-\left(\widetilde{\delta}R\right)^{-\alpha}}{1-R\left(\widetilde{\delta}R\right)^{-\alpha}} \end{aligned}$$

with

$$\widetilde{\delta}^{-\alpha} = e\delta_m^{-\alpha}$$

or

$$\begin{aligned} \widetilde{\delta}^{\alpha} &= \frac{1}{e}\delta_m^{\alpha} \\ &= \frac{a}{1+a}\frac{b+1}{b}\delta_m^{\alpha} \\ &= \frac{\frac{1}{b}+1}{\frac{1}{a}+1}\delta_m^{\alpha} \\ &= \frac{\left(\frac{\mu}{1-\mu}\frac{\delta_f}{\delta_m}\right)^{\alpha}+1}{\left(\frac{\mu}{1-\mu}\right)^{\alpha}+1}\delta_m^{\alpha} \\ &= \frac{\mu^{\alpha}\delta_f^{\alpha}+(1-\mu)^{\alpha}\delta_m^{\alpha}}{\mu^{\alpha}+(1-\mu)^{\alpha}}. \end{aligned}$$

When $s = 0$, the following system results:

$$\begin{aligned} \frac{u'(y-c_{f1})}{u'(c_{f1})} &= \frac{\mu}{1-\mu} \\ \frac{u'(y-c_{f2})}{u'(c_{f2})} &= \frac{\delta_f}{\delta_m}\frac{\mu}{1-\mu} \\ \frac{u'(y-c_{f1})}{u'(y-c_{f2})} &= \delta_m R + \frac{\gamma}{1-\mu}\frac{1}{u'\left(y-c_{f2}\right)} > \delta_m R \end{aligned}$$

Proceeding as above it is not difficult to show that $\gamma > 1$ implies

$$R\widetilde{\delta} < 1.$$

So the household saves a positive amount only when the available return 'beats' the household discount factor. Again this exactly parallels the individual case.

Differentiating the household discount factor with respect to μ Differentiate $\widetilde{\delta}^{\alpha}$ with respect to μ to find:

$$\begin{aligned}
\frac{\partial \widetilde{\delta}^{\alpha}}{\partial \mu} &= \frac{\begin{array}{c}\left[\alpha\mu^{\alpha-1}\delta_f^{\alpha} - \alpha(1-\mu)^{\alpha-1}\delta_m^{\alpha}\right]\left[\mu^{\alpha}+(1-\mu)^{\alpha}\right] \\ -\left[\mu^{\alpha}\delta_f^{\alpha}+(1-\mu)^{\alpha}\delta_m^{\alpha}\right]\left[\alpha\mu^{\alpha-1}-\alpha(1-\mu)^{\alpha-1}\right]\end{array}}{\left[\mu^{\alpha}+(1-\mu)^{\alpha}\right]^2} \\
&= \alpha\frac{\mu^{\alpha-1}(1-\mu)^{\alpha-1}\left(\delta_f^{\alpha}(1-\mu)-\delta_m^{\alpha}\mu+\mu\delta_f^{\alpha}-\delta_m^{\alpha}(1-\mu)\right)}{\left[\mu^{\alpha}+(1-\mu)^{\alpha}\right]^2} \\
&= \frac{\alpha\mu^{\alpha-1}(1-\mu)^{\alpha-1}}{\left[\mu^{\alpha}+(1-\mu)^{\alpha}\right]^2}\left\{\delta_f^{\alpha}-\delta_m^{\alpha}\right\}
\end{aligned}$$

Since $\alpha > 0$, the term in the final curly braces has the same sign as $\delta_f - \delta_m$. Also, $\widetilde{\delta}^{\alpha}$ is an increasing function of $\widetilde{\delta}$. Therefore, $\frac{\partial\widetilde{\delta}}{\partial\mu}$ has the same sign as $\delta_f - \delta_m$.

A.3 Signing $\frac{\partial^2 s^*}{\partial\delta\partial R}$

$$\begin{aligned}
\frac{\partial^2 s^*}{\partial\delta\partial R} &= \frac{\partial}{\partial R}\frac{\partial s^*}{\partial\delta} \\
&= \frac{\partial}{\partial R} y\alpha\delta^{\alpha-1}\frac{R^{1-\alpha}+R^{-\alpha}}{(\delta^{\alpha}+R^{1-\alpha})^2} \\
&\qquad \left((1-\alpha)R^{-\alpha}-\alpha R^{-\alpha-1}\right)\left(\delta^{\alpha}+R^{1-\alpha}\right)^2 \\
&= y\alpha\delta^{\alpha-1}\frac{-2\left(\delta^{\alpha}+R^{1-\alpha}\right)(1-\alpha)R^{-\alpha}\left(R^{1-\alpha}+R^{-\alpha}\right)}{(\delta^{\alpha}+R^{1-\alpha})^4} \\
&= y\alpha\delta^{\alpha-1}\frac{\left((1-\alpha)R^{-\alpha}-\alpha R^{-\alpha-1}\right)\left(\delta^{\alpha}+R^{1-\alpha}\right)-2(1-\alpha)R^{-\alpha}\left(R^{1-\alpha}+R^{-\alpha}\right)}{(\delta^{\alpha}+R^{1-\alpha})^3} \\
&= \frac{y\alpha\delta^{\alpha-1}R^{-\alpha}}{(\delta^{\alpha}+R^{1-\alpha})^3}\left\{\left(1-\alpha-\alpha R^{-1}\right)\left(\delta^{\alpha}+R^{1-\alpha}\right)-2(1-\alpha)(R+1)R^{-\alpha}\right\} \\
&= \frac{y\alpha\delta^{\alpha-1}R^{-\alpha}}{(\delta^{\alpha}+R^{1-\alpha})^3}\left\{R^{-\alpha}\left[R(1-\alpha)-\alpha-2(1-\alpha)(R+1)\right]+\delta^{\alpha}\left(1-\alpha-\alpha R^{-1}\right)\right\} \\
&= \frac{y\alpha\delta^{\alpha-1}R^{-\alpha}}{(\delta^{\alpha}+R^{1-\alpha})^3}\left\{R^{-\alpha}\left[R(\alpha-1)+\alpha-2\right]+\delta^{\alpha}\left(1-\alpha-\alpha R^{-1}\right)\right\}
\end{aligned}$$

This is positive whenever the expression in curly braces is positive, that is when

$$\begin{aligned}
R^{-\alpha}\left[R(\alpha-1)+\alpha-2\right]+\delta^{\alpha}\left(1-\alpha-\alpha R^{-1}\right) &> 0 \\
R^{-\alpha}\left[R(\alpha-1)+\alpha-2\right] &> \frac{\delta^{\alpha}}{R}\left(R(\alpha-1)+\alpha\right) \\
R(R\delta)^{-\alpha} &> \frac{R(\alpha-1)+\alpha}{R(\alpha-1)+\alpha-2} \\
R(R\delta)^{-\alpha} &> 1+\frac{2}{R(\alpha-1)+\alpha-2} \\
R(R\delta)^{-\alpha}-1 &> \frac{2}{(R+1)(\alpha-1)-1}.
\end{aligned}$$

Though it is generally not determined whether this inequality holds, it *does* hold when $\alpha < 1$: Conditional on $\alpha < 1$ the left-hand side is strictly positive since

$$R > (R\delta)^{\alpha},$$

and the right-hand side is strictly negative.

So the sign of $\frac{\partial^2 s^*}{\partial \delta \partial R}$ is positive when $\alpha < 1$ but generally indeterminate.

A.4 Signing $\frac{\partial^2 s^*}{\partial \delta^2}$

$$\begin{aligned}
\frac{\partial^2 s^*}{\partial \delta^2} &= \frac{\partial}{\partial \delta} y \frac{\alpha \delta^{\alpha-1} R^{-\alpha}(R+1)}{\left(\delta^{\alpha}+R^{1-\alpha}\right)^2} \\
&= y\alpha R^{-\alpha}(R+1)\frac{(\alpha-1)\delta^{\alpha-2}\left(\delta^{\alpha}+R^{1-\alpha}\right)^2 - \delta^{\alpha-1}2\left(\delta^{\alpha}+R^{1-\alpha}\right)\alpha\delta^{\alpha-1}}{\left(\delta^{\alpha}+R^{1-\alpha}\right)^4} \\
&= \frac{y\alpha R^{-\alpha}(R+1)\delta^{\alpha-2}}{\left(\delta^{\alpha}+R^{1-\alpha}\right)^3}\left\{(\alpha-1)\left(\delta^{\alpha}+R^{1-\alpha}\right) - 2\alpha\delta^{\alpha}\right\} \\
&= \frac{y\alpha R^{-\alpha}(R+1)\delta^{\alpha-2}}{\left(\delta^{\alpha}+R^{1-\alpha}\right)^3}\left\{R^{1-\alpha}(\alpha-1) - \delta^{\alpha}(\alpha+1)\right\}
\end{aligned}$$

which is negative when $\alpha < 1$ but generally indeterminate.

2 Funeral insurance

2.1 Introduction

Funeral insurance is probably one of the earliest forms of insurance. Its popularity is well documented throughout history and across the globe, and it remains the single most popular insurance type in large parts of today's Africa. Yet the study of this phenomenon has been largely neglected by economists. The purpose of this paper is to provide an explicit treatment of funeral insurance as a distinct form of insurance.

Funeral insurance is not life insurance. What these two types of insurance have in common is that the covered event is the death of a specified individual. But it is a remarkably consistent feature of death-triggered insurance, both historically and in modern-day developing countries, that the payout takes the form of funeral-related goods and services. In contrast, a life insurance policy typically pays out in cash. It is a central claim of this paper that the distinction is fundamental, and that the type of insurance preferred will depend on the circumstances and characteristics of the decision-maker.

Today insurance companies in many parts of the world offer formal funeral insurance policies. But the history of funeral insurance, and by implication the history of insurance more generally, is closely linked to that of funeral societies. From ancient Greece and Rome, via mediaeval Europe and Victorian Britain, to large parts of modern-day Sub-Saharan Africa, the most common way to take out funeral insurance has been to join a funeral society. The primary function of these groups is precisely to pool the risk associated with the death of members or their close relatives, by using members' contributions to organise funerals. Historians have argued that funeral societies are the precursors of modern insurance companies (Fingland Jack 1912, Trenerry 1926). More recently, funeral societies have attracted the attention of economists as instances of informal or semi-formal risk-pooling groups.

South Africa may be unique in that both formal and informal funeral insurance, as well as life insurance, are widely held. Traditional funeral societies compete with modern insurance companies in providing funeral insurance, and the latter also offer a range of life insurance policies. Traditional African communities coexist with a fully industrialised modern economy and world-class formal-sector financial institutions. South Africa is therefore close to being the ideal testbed for the ideas presented here.

The analysis of funeral insurance as a distinct form of insurance is the main contribution of this paper. Its ancient roots and importance in the contemporary developing world are highlighted. It is also emphasised that even though the association between funeral insurance and funeral societies has been and continues to be important, the concept of funeral insurance should be distinguished from its implementation.

A model is presented in which funeral insurance is contrasted with life insurance. Whereas previous work on informal insurance has focused on the sustainability of the

contract, this paper abstracts from organisational form and asks under what circumstances funeral insurance will be preferred to general life insurance. In order to focus on the mechanism of interest, none of the two insurance types are presented with a cost or efficiency advantage over the other. Funeral insurance is modelled as life insurance bundled with a constraint on the use of the payout. Therefore, funeral insurance can serve as an inter-generational commitment device when the preferences of the policyholder (parent) differs from those of the beneficiary (child). Conditions are derived under which each type of insurance is preferred: For low levels of wealth, and for high levels of income, the agent will take out funeral insurance. For intermediate levels of wealth and income she prefers life insurance, or a mix of life and funeral insurance at a non-binding level. For high levels of wealth and low income, the agent prefers not to take out any insurance. The model also contributes to the insurance literature by suggesting that even conventional life insurance is not necessarily motivated by altruism in the usual sense of the word.

The theoretical findings are held up against reality in the form of empirical analysis of data from South Africa. A large, nation-wide marketing survey usually available only in summary form is used here for the first time to run robust regressions of insurance choices on personal and household characteristics. The results of the analysis are consistent with one of the predictions of the theoretical model, but data limitations permit only a very limited empirical test.

The rest of the paper is organised as follows. The next section provides an introduction to funeral societies and funeral insurance. The following section describes the importance of funeral and life insurance in contemporary South Africa. Thereafter a brief overview of the relevant literature is given, before the model is presented, solved and interpreted. Next is the empirical section, and finally a conclusion.

2.2 Funeral societies and the history of insurance

Funeral societies are mutual risk-pooling groups designed to ensure decent funerals for members and/or other persons nominated by them, typically close relatives or household members. On the death of a covered person, the group will provide some combination of cash, labour and goods towards the funeral. Many collect fixed cash premia at regular intervals, while some transact only when a death occurs.[11]

[11] Funeral societies should be distinguished from rotating and accumulating savings and credit associations (ROSCAs/ASCAs). In all these types of arrangement, a fixed amount of money is typically collected regularly from each member. But in a ROSCA payout rotates amongst the members in a systematic fashion, typically determined either randomly or by bidding. During a cycle of ROSCA operation, each member will receive the payout exactly once (Besley, Coate & Loury 1993). ASCAs operate on a similar principle but, as the name suggests, funds may be accumulated and held or invested jointly rather than being paid out immediately after collection. In contrast to these arrangements, funeral society payouts are triggered by random events (the death of a covered person) and there is no guaranteed relationship between contribution and payout over the the period of membership. Therefore ROSCAs and ASCAs are fundamentally savings and credit devices (though various other functions have been suggested in the literature), whereas a funeral society is at the core a risk-pooling arrangement. Perhaps MIA - mutual insurance association - would be a suitable acronym?

Funeral societies have a long history and global reach. Solon the Athenian statesman (ca. 638–558 BC) passed a law regulating their activity (Parrott 1985). They were widespread during the Roman empire, operating on basic principles that are virtually indistinguishable from those found in many parts of the world today.[12] In mediaeval Europe they were organised by professional guilds. During the industrial revolution in England, funeral societies could be set up as local community groups, or organised as large friendly societies (Cordery 2003). For a vivid description of the importance of burial societies in British working class life in the late 19th and early 20th centuries, see Johnson (1985).

The history of funeral societies is yet to be written, but Van der Linden (1996) covers the history of mutual benefit societies more generally in 26 countries across Europe, North and South America and Asia. Though many of the institutions described also cover events other than death, funeral cover seems to have been an important component of most of them, and not uncommonly the only one. No doubt their enduring popularity is in part due to the inherent insurability of death risk: Moral hazard is generally less of a problem than for other common risks.

Pure funeral societies still exist in rich countries, though they are now of little importance compared to formal insurance (public or private). But funeral societies are still popular in the developing world, particularly in Sub-Saharan Africa.[13]

Funeral societies are often said to be informal insurance providers, but it is worth clarifying what is meant by 'informal' in this context. As Dercon et al. (2006) point out, funeral societies often operate on a clearly defined, sometimes written, set of rules. The terms of the 'policy' are known in detail, including who is covered, conditions of cover and size of premia and payouts. Often there is also a system of fines for non-compliance with these rules, and at least in South Africa it is not uncommon for funeral societies to have a special uniform which is compulsorily worn by members at meetings and funerals. Strict rules are the norm rather than an exception. However, most funeral societies are not registered with the authorities and hence not regulated. They are not part of the de-personalised formal economy. In particular, it is doubtful whether a member of a funeral society who feels unjustly treated has recourse to

[12] "Early in the Empire, associations were formed for the purpose of meeting the funeral expenses of their members, whether the remains were to be buried or cremated, or for the purpose of building columbāria, or for both. These co-operative associations (collēgia fūnerātīcia) started originally among members of the same guild [..] or among persons of the same occupation. They called themselves by many names, cultōrēs of this deity or that, collēgia salūtāria, collēgia iuvenum, etc., but their objects and methods were practically the same. If the members had provided places for the disposal of their bodies after death, they now provided for the necessary funeral expenses by paying into the common fund weekly a small fixed sum, easily within the reach of the poorest of them. When a member died, a stated sum was drawn from the treasury for his funeral, a committee saw that the rites were decently performed, and at the proper seasons [..] the society made corporate offerings to the dead." From Johnston (1903).

[13] A non-exhaustive literature search identified academic work with references to burial societies in contemporary Botswana (Ngwenya 2003), Zimbabwe (Dhemba, Gumbo & Nyamusara 2002), Ethiopia and Tanzania (Dercon et al. 2006), Ghana (de Witte 2003), India (Rutherford 2000) and Thailand (Bryant & Prohmmo 2002) as well as South Africa.

the enforcement apparatus of the formal judiciary. This clearly distinguishes funeral society membership from policies offered by formal insurance companies.[14]

However, funeral insurance is not provided only by funeral societies. A quick internet search for 'funeral insurance' at the time of writing this was enough to confirm that formal insurance companies currently offer specific funeral policies in several rich countries including the USA and the UK.

One might object that funeral societies offer more than insurance. After all, whereas formal insurance is 'only' a financial product, funeral society colleagues are often also friends who may provide a sense of belonging and support, especially in times of grief. Nonetheless, virtually all writers on the subject agree that participants think of their funeral societies as a financial arrangement first and foremost. In support of this, a nationally representative South African survey asked funeral society members why they belonged to these groups. The three most popular responses were "Help with funeral arrangements" (79%), "To help when there is a death" (53%), and "Provide funerals the family deserves" (24%). Many fewer selected the response "To provide comfort and support" (13%) or "To socialise" (4%) (FinMark Trust 2003). Likewise, Dercon et al. (2006) report that in both Ethiopia and Tanzania, the primary focus of funeral societies is to provide funeral cover. It is undeniable that funeral societies have social aspects, but these alone cannot explain the popularity of the institution.

The history of funeral insurance is closely linked with that of funeral societies, and even today funeral insurance is much more likely than life insurance to be offered informally. Nevertheless, formal-sector funeral insurance exists, and so do informal insurance groups which are not exclusively concerned with funeral cover. It is a key message of this paper that the concept of funeral insurance should be distinguished from the organisational form of the provider. Neither depends on the other, though funeral insurance appears to be popular in circumstances that favour informal provision.

2.3 Funerals and funeral insurance in South Africa

Funerals are expensive events in Africa[15], and South Africa is no exception. Case, Garrib, Menendez & Olgiati (2008) find that funerals amongst black households in a region of KwaZulu-Natal on average costs a median annual income. Roth (2001) conducted in-depth interviews with twelve households in a township in Eastern Cape, ten of whom had recently been involved in the funeral of a relative. He reports that the average expenditure on a funeral in his small sample was 15 times monthly household income.

[14] It is possible that Bloch, Genicot & Ray (2008) have in mind another understanding of 'informal' when they write that most informal insurance is bilateral and that "our existing idea of insurance as taking place within an explicit "club" of several people may be misleading".

[15] Arhin (1994) describes lavish funeral customs in Ghana, and Adamolekun (2001) does the same for Western Nigeria. Dercon et al. (2006) report "substantial" costs in Ethiopia and Tanzania.

The high cost of funerals are related to their great cultural importance in Sub-Saharan Africa. Traditional belief holds that the spirits of the dead can influence the lives of the living and must be treated with respect. Often relatives will come from far away and the cost of catering for them for several days can be substantial (Roth 2001). But often there is also a strong component of 'conspicuous consumption', and a sense of shame if the deceased's family cannot afford an elaborate funeral.[16] Case et al. (2008) deem the importance of funerals to be ahead of births, graduations and weddings in traditional family and community life.

The social importance of funerals, and no doubt also the death toll of the AIDS epidemic, partly explain why both formal and informal funeral insurance are 'big' in South Africa. Porteous & Hazelhurst (2004) conclude that funeral society membership is around 18% of the population (8 million out of 45 million). As a proportion of black adults, membership is even higher. In the formal sector, funeral insurance is the single most popular type of policy, with about 8% (3.5 million) of the population being policy-holders. Since a typical policy will cover close relatives as well as the policy-holder, the proportion of the population that is covered by some form of funeral insurance is substantially larger than these numbers suggest.

2.4 Related literature

This paper relates to earlier work on the demand for life insurance. Modern economic analysis of the problem starts with Yaari (1965), who introduces life insurance as a way of coping with uncertain lifetime in a model with either a bequest motive or a credit market combined with a non-negative terminal wealth requirement. Fischer (1973) characterises insurance demand functions in a discrete-time model. He finds that an agent who lives off his wealth is unlikely to ever purchase insurance, a result which will have a parallel in this paper. Lewis (1989) was the first to consider the point of view of the insurance beneficiaries, a consideration which will be important in this paper. He lets insurance demand be determined by the child beneficiary, on whose behalf life insurance is arguably purchased.

Economists have long studied the question of why people leave bequests. Some believe that bequests are accidental and caused by the combination of uncertain lifetimes and imperfect insurance markets. On the other hand, Bernheim (1991) uses evidence from the demand for life insurance to argue that bequests are intentional. This paper nuances the discussion by positing that even if they are intentional, bequests are not necessarily altruistic in the usual sense. Demand for insurance may be motivated by a concern for specific types of post-mortem expenditure that have more to do with preserving one's own good name and afterlife than with 'warm glow' or an empathetic concern with the welfare of the next generation.

[16] In some funerals the price tag on the coffin is not removed, but displayed for everyone to see. Warnecke (1994) as quoted by Thomson & Posel (2002).

In contrast to this paper's focus on a specific type of insured event, much of the earlier work on risk-coping in developing countries have looked at consumption smoothing in general (Besley 1995*a*, Morduch 1999, Dercon 2002). Townsend (1994) tests for, and rejects, full risk-sharing in Indian villages. This finding has been verified by later studies, and the robustness of the result has inspired a substantial body of theoretical work relating imperfect enforcement to bounds on the risk-sharing contracts that may be written (Ligon, Thomas & Worrall 2002).

Arnott & Stiglitz (1991) is perhaps the best-known article on the interaction of formal and informal insurance. The authors find that in the presence of moral hazard and formal insurance, informal insurance is beneficial if the informal insurers have an information advantage, but can be harmful if not. This paper abstracts completely from moral hazard, arguing that death-triggered insurance is much less likely to be subject to this problem than many other types of insurance. This is in line with Fafchamps & Lund (2003) who find that funerals are better insured with informal gifts and loans than other events such as crop failure and mild illness.

A small literature focuses on commitment devices in developing countries, a good example of which is Ashraf, Karlan & Yin (2006*a*). These are related to the theoretical literature on hyperbolic discounting (Harris & Laibson 2003). But not much is known about devices that operate between generations.

This paper parallels parts of the literature on a more well-known (to economists) form of group-based informal finance: The rotating savings and credit association (ROSCA). Levenson & Besley (1996) look at the determinants of ROSCA participation in Taiwan. In a situation semblant of that of funeral societies in South Africa, they deem ROSCA membership in Thailand surprisingly high (at least a fifth of all households are members) for a newly industrialised country.

Funeral societies have received relatively little attention from economists. Dercon et al. (2006) discuss funeral societies in Ethiopia and Tanzania, and Bold (2007) test for risk-sharing in Ethiopian societies. Bryant & Prohmmo (2002) ask why burial society premia in a village in North-Eastern Thailand are equal for all households irrespective of risk. A good source of information on the workings and characteristics of funeral societies in South Africa is Thomson & Posel (2002). Roth (2001) is a case study of formal and informal funeral insurance in a rural South African township. Ardington & Leibbrandt (2004) find a strong correlation between formal employment and take-up of funeral insurance. Keswell (2004) looks at the relationship between employment and membership of informal networks such as ROSCAs and funeral societies.

2.5 Theory

2.5.1 The model

There are two agents, a parent and a child, and two periods. The parent faces a risk of death. In the first period she allocates her starting wealth w between own

consumption, insurance and savings. With probability q she survives to the second period, receives an income y and consumes all her resources before the game ends. But with probability $1-q$ she dies at the end of the first period. If so, the child inherits her savings and receives any insurance payout. The child allocates her resources between the parent's funeral and her own consumption before the game ends.

The parent's expected utility at the beginning of the first period is

$$V = u(C_1) + \beta \{qu(C_2) + (1-q)\phi(f)\},$$

where C_1, C_2 are the parent's consumption variables, f is funeral expenditure in the case of death, and $u()$ and $\phi()$ are the instantaneous utility functions for consumption and the funeral, respectively. β is a discount factor.

If the parent dies at the end of the first period, the child is the active agent in the second period. She derives utility from the funeral as well as from her own consumption,

$$U = \gamma\phi(f) + (1-\gamma)u(c-a).$$

Here, $\gamma \leq 1$ is a weight which determines how the child values funeral expenditure relative to her own consumption, c. The constant $a \geq 0$ has the effect of increasing the emphasis on consumption for lower resource levels. It may be thought of as defining a minimum consumption level: As the child's available resources decrease towards a, the child will prioritise consumption to the detriment of expenditure on the parent's funeral.

Note that in this set-up there is no altruism in the usual sense of the word. However, both agents derive utility from the parent's funeral. This may be thought of as a common concern for a public good: The family's reputation or 'name'. It may also be thought of as the parent's utility from the anticipation of a dignified departure, and for the child it may form part of the grief process.

There are two types of insurance, both of which pay out in the event that the parent dies. A life insurance payout can be disposed of as the child wishes, whereas a funeral insurance payout can only be spent on the parent's funeral. This is the only difference between the two types. Both types of insurance are equally priced with premia equal to $\alpha > 1$ times the actuarially fair rate.

2.5.2 A solution

In order to arrive at an analytical solution, some further simplifying assumptions are made: The parent's consumption in the first period is fixed. Then w can be renormalised to be net of first-period consumption, and the utility associated with it may be disregarded. Notation may also be simplified by writing $C \equiv C_2$. Time discounting is no longer important and may be ignored ($\beta = 1$), and it is also assumed that there are zero returns on savings.

Also, specify

$$u(x) = \phi(x) = \ln(x).$$

The logarithmic utility function has the advantage of being tractable, and also satisfies constant relative risk aversion (CRRA).

The parameters w, y and a are taken to be non-negative, and q and γ must be contained in the interval $(0,1)$. Actuarially unfair insurance requires $\alpha > 1$. These parameter assumptions are quite innocuous, but a final assumption is more substantial:

$$\alpha(1-q) < \frac{1+q-\frac{q}{\gamma}}{1+q}. \tag{2.1}$$

Intuitively, insurance cannot be *too* unfair. The robustness of the theoretical findings with respect to relaxation of this constraint will be discussed below. Note that right-hand side is smaller than unity, so the inequality implies the natural requirement that the premium should not exceed the payout.

The assumptions required for an analytical solution to the problem are now in place. The child's problem is to maximise

$$U = \gamma \ln(f) + (1-\gamma)\ln(c-a)$$

subject to her budget constraint

$$f + c = W,$$

where W is the child's total resources, and

$$f > F$$

where F is the funeral insurance payout (which may be zero.) The solution defines $f(W,F)$; funeral expenditure as a function of child resources and funeral insurance cover.

The parent's problem is to maximise her expected utility

$$V = q\ln(C) + (1-q)\ln(f(W,F)),$$

with respect to life and funeral cover L and F, subject to the constraints

$$\begin{aligned} C &= w_s \\ W &= w_d \\ w_s &= w - \alpha(1-q)(L+F) + y \\ w_d &= w - \alpha(1-q)(L+F) + (L+F) \\ L &\geq 0 \\ F &\geq 0. \end{aligned}$$

Here, w_s and w_d are the total available resources at the beginning of the second period in the cases of survival and death, respectively. In both cases, the amount saved from the first period is just the starting wealth minus insurance premium. In the survival case the parent receives the income y in addition, while in the death case the child receives any insurance payout. Negative insurance cover is ruled out.

The primary interest is in the parent's decisions with regard to insurance cover, L and F. The solution is found by applying the standard method of backward induction: The child's solution is identified first, conditional on the parent's decisions. The child's solution is substituted into the parent's insurance problem, which can now be solved as a single-period problem.

The following proposition characterises the parent's optimal insurance decisions, L and F.

Proposition: There are three cases, depending on the magnitude of y relative to w.

1. *Binding funeral insurance.* For

$$y > \frac{\kappa w - a}{1 - \kappa}$$

with

$$\kappa = \frac{\gamma - (\gamma (1 - \alpha (1 - q)))^q}{\gamma \alpha (1 - q)},$$

the solution is

$$\begin{aligned} F &= \frac{w + y}{\alpha} \\ L &= 0, \end{aligned}$$

and the funeral insurance payout is binding on funeral expenditure.

2. *Non-binding insurance.* For

$$\frac{(\alpha - 1) w - \alpha q a}{1 - \alpha (1 - q)} < y < \frac{\kappa w - a}{1 - \kappa},$$

total insurance cover is

$$F + L = \frac{(1 - \alpha (1 - q)) y - (\alpha - 1) w + \alpha q a}{\alpha (1 - \alpha (1 - q))}.$$

Moreover, the breakdown of cover into L and F is indeterminate except that F must satisfy

$$F < \gamma \left(\frac{1 + \alpha q}{\alpha} w + \frac{1 - \alpha (1 - q)}{\alpha} y - a \right)$$

so that it is not binding on the child's funeral expenditure.

3. *No insurance.* For

$$y < \frac{(\alpha - 1) w - \alpha q a}{1 - \alpha (1 - q)},$$

the parent does not take out any insurance.

Proof. The proof is in the appendix.

Note that as long as funeral insurance is not binding on the child's funeral expenditure, it is equivalent to life insurance. This must be so since the only difference between the two types of insurance is that funeral insurance payouts must be spent on the funeral. Hence in case 2, the parent is indifferent between a pure life policy, and any mix of life and funeral cover which sums to the same amount, as long as the funeral insurance constraint is not binding on the child. This case may therefore be called 'non-binding insurance'.

The result is illustrated in Figure 2.1 and may be characterised as follows. Hold y fixed at an arbitrary level. For sufficiently low w, the parent takes out funeral insurance only. The level of cover is binding on the child's expenditure decision. As w increases, an increasing amount of funeral cover is taken out until a boundary is crossed at which there is a discontinuity in the amount of insurance held and beyond which insurance is not binding. As w continues to increase, total cover decreases but the parent is indifferent to its breakdown in F and L as long as F stays below the level at which it binds the child. At some point insurance cover reaches zero. As w increases beyond this point the parent does not take out any insurance.

Intuitively, the parent has three options. The first is to incur the unfair insurance premium on the whole funeral expenditure amount by purchasing a binding amount of funeral insurance. The second is to 'top up' the child's inheritance with non-binding insurance. In this case only a fraction γ of the total inheritance will go towards funeral expenses, but the advantage is that the actuarially unfair insurance mark-up α is only paid on the top-up as opposed to on all of the funeral expenditure amount. Finally, there is the option of not taking out any insurance, in which case the child will simply spend a fraction γ of inherited wealth on the funeral.

At low levels of inheritable wealth w relative to survival-state income y, funeral expenditure in the absence of insurance is low relative to y. Therefore, the parent wishes to transfer resources from the survival-state to the death-state by taking out binding funeral insurance. As inheritable wealth increases relative to y, the marginal benefit of transferring resources from the survival state to the death state decreases, as does the willingness to pay the unfair mark-up on the insurance premium. At some point, the value of 'topping up' the inheritance with a smaller amount of non-binding insurance becomes preferable, because the advantage of not paying the actuarially unfair mark-up on the whole funeral expenditure outweighs the loss associated with the fact that the child spends only a fraction γ of her resources on the funeral. Beyond this wealth level, the parent takes out non-binding insurance. As wealth increases further, there is a point beyond which the cost of the actuarially unfair insurance outweighs the benefit of consumption smoothing altogether. Beyond this point, the parent does not take out any insurance.

The least innocuous assumption underlying the above proposition is arguably (2.1).

Intuitively it says that the insurance premium cannot be 'too unfair'. Though this paper will proceed on the basis that the assumption holds, it is interesting to consider how the parent's decisions would change if it were violated. A formal result will not provided, but it can be shown that the qualitative effect of violating the assumption is to eliminate case 2 of the proposition, namely the 'non-binding' case. Thus there would be only two cases, binding funeral insurance and no insurance. Effectively, there would be zero demand for life insurance.

In non-mathematical terms, the reason is as follows: Whether (2.1) holds or not, binding funeral insurance is preferred when wealth is sufficiently small relative to income. As wealth increases while income is held fixed, the gain from insurance diminishes, because wealth is available in both the survival and the death states. This holds for both life (or more generally, non-binding) and funeral insurance. There are three critical level as wealth increases: a) One beyond which non-binding is preferred to binding insurance, b) one beyond which no insurance is preferred to binding funeral insurance, and c) one beyond which no insurance is preferred to non-binding insurance. Assumption (2.1) determines the ordering of these critical levels. As long as the assumption holds, threshold a) is lower than b) and c). This gives rise to a region of non-binding insurance between a) and c). If (2.1) does not hold, then c) is lower than a), implying that the parent will transit directly from binding funeral insurance to no insurance as wealth increases.

2.5.3 Mapping theory to data

The aim of the next section is to hold the theoretical findings up against the light in the form of empirical analysis of household data. In order to do this it is necessary to extract from the theory some predictions regarding the relationships of quantities that are observable in the data.

The main predictions from the model relate wealth and income to the insurance decisions taken by the parent. As will become clear below, the data makes it difficult to control for wealth with any precision. The analysis will therefore largely be limited to analysing how insurance choice depends on certain crude measures of wealth, when income is held fixed. The main testable prediction is that increasing wealth for a fixed level of income is associated with a transition from funeral insurance to life insurance. Increasing wealth further is expected to be associated with a transition out of insurance.

Since there is no information on insured amounts, premia or other contractual terms in the data, the focus of the empirical analysis will be on the binary take-up decisions of the household with respect to funeral and life insurance. As the insurance decisions are likely to be jointly determined, a natural choice of empirical specification is the bivariate probit:

If $\widetilde{L}$ and $\widetilde{F}$ are the binary decision variables, then let

$$\begin{aligned} \widetilde{L} &= \begin{cases} 1, & \text{if } X\beta_L + \varepsilon_L > 0 \\ 0, & \text{otherwise} \end{cases} \\ \widetilde{F} &= \begin{cases} 1, & \text{if } X\beta_F + \varepsilon_F > 0 \\ 0, & \text{otherwise} \end{cases} \end{aligned} \tag{2.2}$$

Here, X is a matrix of explanatory variables common to both equations. It will include measures of wealth and income along with control variables. The error terms are bivariate normally distributed with expectation 0, variance 1 and covariance ρ,

$$\begin{pmatrix} \varepsilon_L \\ \varepsilon_F \end{pmatrix} | X = \mathcal{N}\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix} \right). \tag{2.3}$$

Given this specification and data, the parameters β_L, β_F and ρ can be estimated using maximum likelihood. The primary interest is in those components of β_L and β_F which relate to wealth and income.

Apart from variables relating to income, wealth and insurance cover, the model also has parameters for the probability of survival q, the weight of funeral expenditure in the child's welfare function γ, minimum consumption a, and insurance efficiency parameter α. These are not easily translated into observable quantities in the data available, and will therefore be captured in the error terms.

2.6 Empirics

2.6.1 Data

The data are taken from the 2004 wave of the *All Media and Products Survey* published by the South African Advertising Research Foundation. It is part of a long-running series of surveys covering South Africans aged 16 and over. The sample of 24,489 overrepresents the higher income brackets, but weights inversely proportional to the probability of selection are provided to allow inferences about the underlying population. The weights are used throughout the following analysis. The survey's emphasis is on commercial product usage, but there are also sections covering personal data, family and housing situation, and leisure activities.

However, the survey is primarily an individual, rather than a household, survey. The survey respondent is a randomly chosen person above the age of 16 in the selected household. Most of the questions relate to the survey respondent and there is limited household-level information. A further problem is that there is only binary information for a number of the most relevant variables. For instance, respondents are asked whether they have a life cover policy, but no further details on the policy are provided. The data appendix provides more detail on the construction of the variables used.

Total monthly household income data is provided as a discrete variable, i.e. as

falling in one of 33 income ranges. The income variable used here is the logarithm of the midpoint of the reported income range.

The main theoretical predictions are the relationships between household wealth and income and insurance take-up. However, the available data limits the extent to which the predictions may be tested. The ideal would be to have sufficiently rich data to look at the effect of changing wealth while holding income constant, and vice versa. This is all the more important since wealth and income are typically strongly correlated. However, there is no precise measure of wealth in the data. In the following analysis, car and home ownership are used to proxy for wealth. This is fine as far as studying the predictions of wealth-holding is concerned, because it is possible to control for income at a more fine-grained level. But it makes it difficult to isolate the importance of income on insurance decisions, because it is not possible to control for wealth holdings to any level of accuracy.

Weighted summary statistics are provided in Table 2.1. They suggest that 22% of South Africans aged 16 or over belong to a funeral society. This is somewhat lower than earlier estimates (it corresponds to 6.6 million members versus 8 million as suggested by Porteous & Hazelhurst (2004)). Furthermore, 12% personally have formal funeral insurance. Given the data it is not possible to distinguish those who do not personally have any formal or informal funeral insurance but is covered by someone else's policy, from those who are not covered at all. 11% of the population have life insurance.

2.6.2 Main results

As a graphical precursor to the results, consider Figure 2.2. Here the take-up of funeral insurance and life insurance are plotted against household income. The general pattern appears to be in line with the theoretical prediction that the take-up of funeral insurance is not very sensitive to income. On the other hand, life insurance take-up shows a strong positive relationship with income. But the chart does not hold wealth fixed. One could imagine a similar chart where total assets or wealth is on the horizontal axis, but scalar data on wealth is not available from the survey.

The empirical specification is given by (2.2),(2.3). The control variables which will be used are binary variables coding respondent age in decadal bands, rural location, province, whether the respondent has own children and whether he/she is married/cohabiting. The observation weights are used in all regressions, and standard errors are robust and clustered at the province level.

The main results are presented in Table 2.2. Each bivariate probit regression is reported over two columns, one for funeral insurance and another for life insurance. In regression 1 the take-up variables are regressed on wealth with income as the only control variable. For funeral insurance the coefficients on the wealth variables are negative but not significantly different from zero. For life insurance the coefficients

are positive and significant (but only marginally for home ownership). In regression 2 the controls for age, rural location and province are included. For funeral insurance the wealth coefficients are still negative and now significant (but only marginally so in the case of home ownership). Income remains insignificant, while car ownership is negative and significant. For life insurance the coefficient on car ownership remains positive and highly significant, whereas the coefficient on home ownership is not. Regression 3 adds controls for having own children and being married or cohabiting. Though potentially endogenous, the latter two controls are interesting because insurance behaviour is likely to change over the lifecycle. The results are qualitatively similar to those for regression 2.

Overall then, it seems that owning a car makes funeral insurance less likely and life insurance more likely, holding income fixed. This is in line with the predictions. The results for home ownership are similar in magnitudes, but the coefficients are mostly not significant or only marginally significant.

It is worth pointing out the potential hazard in putting household income and indicators of wealth on the right-hand side of all these equations. Clearly there are concerns of endogeneity and a real possibility of joint determination of income and wealth on the one hand, and insurance decisions on the other. For example, both might be related to variation in unobservable ability which is captured in the error terms. But it is difficult to avoid this problem given that the main predictions of the theoretical model concern precisely the relationship between insurance, income and wealth. The upshot is that all findings presented here must be interpreted with caution: They are indicative, rather than watertight.

These regressions were run on the sample as whole. But ethnic groups are still largely segregated in South Africa, and arguably have different cultures. It is also true that funeral societies are primarily associated with black (African) and coloured people, rather than with the white and Indian sub-populations. As both wealth and cultural affinity may be predicated on ethnicity, the above regressions may in fact be confounding cultural differences with wealth effects. To investigate this, the analysis is repeated for the subsample of black and coloured households alone.

The results are presented in Table 2.3. As above, and in line with theory, car ownership is found to be consistently predictive of life insurance. However, though the relation between home ownership and life insurance is positive in magnitude, it is not significant except in regression 1. In all three regressions, none of the coefficients relating car and home ownership to funeral insurance are significant at the 5% level. Since all significant coefficients are in line with the predictions, it is probably fair to say that these results are at least weakly supportive of the theory.

2.6.3 Funeral insurance and the informal sector

Hitherto the focus has been on funeral insurance as a distinct form of insurance, rather than on its implementation. However, no treatment of funeral insurance would be complete without at least a brief discussion of its close ties with the informal sector.

The availability of formal-sector funeral insurance, along with the existence of informal and semi-formal risk pooling groups that do not provide funeral insurance, provide evidence that the link between funeral insurance and the informal sector is not one of logical necessity or sufficiency. Nevertheless, a large proportion of the people who have, or have ever had, funeral cover obtained it by joining a funeral society. Likewise it seems that funeral insurance has been a particular preoccupation of informal and semi-formal insurance groups throughout history and remains so today.

It is natural to ask what caused these correlations of funeral insurance with the informal sector on the one hand, and of life insurance with the formal sector on the other. It could be that the societal parameters which favour funeral insurance are the same as those that favour informal sector solutions. Or there could be something inherent in funeral insurance which makes it more widespread in economies dominated by the informal sector.

It is often posited that informal sector finance may have an information and/or enforcement advantage over the formal sector. However, it is not immediately clear how these would favour funeral insurance over life insurance. It is also true that informal risk pooling groups may be able to offer moral support, comfort and friendship where formal insurers do not. But again it is not obvious how these factors would favour funeral insurance over life insurance.

Using the model above, funeral insurance can be associated with low wealth relative to income, low survival probability, and a low weight on the parent's funeral in the child's payoff function. To the extent that these factors are also associated with the informal sector, they may partly explain why funeral insurance is so often provided informally.

Though important, these matters will not be further discussed here. There is, however, one aspect of the link between funeral insurance and the informal sector which has a direct bearing on the empirical findings presented above: Since funeral insurance is primarily provided by the informal sector (read: funeral societies) even in South Africa, there is a possibility of confounding the correlates of funeral insurance with those of the informal sector more generally. These concerns may be partially alleviated by ignoring informal insurance and comparing formal funeral insurance with formal life insurance. That is the aim of the next set of regressions.

The results are presented in Table 2.4 . The only estimate which is consistently distinguishable from zero is the positive relation between car ownership and life insurance, and this aligns well with the theory. Regression 1 finds a positive and significant relation between home ownership and funeral insurance, which runs counter to the

predictions. However, when further controls are added in regressions 2 and 3 the coefficient is small, negative and not significant. Like in full-sample analysis, home ownership is positively related life insurance, but the coefficient is only significant in regression 1. Again, the findings are at least weakly supportive of the theoretical predictions.

2.6.4 Further analysis

One concern with the findings so far is related to the fact that the survey was primarily an individual, rather than a household, survey. Some information, such as the income and asset ownership variables used here, was nonetheless collected at the household level. But importantly, the insurance take-up variables are specific to the respondent. In other words, the data tell us whether the respondent has *personally* taken out funeral or life insurance, but not whether the household as a whole has insurance. Some households that have insurance will not appear as such in the data because the respondent is not the main policy-holder. It is likely that the household head or his spouse are the policy-holders in most cases.

One way to counter this problem would be to restrict the analysis to the subsample of households where the respondent is likely to be a financial decision-maker, but it is not clear that the identification of decision-makers could be done with enough precision to be worthwhile. However, another option is to look at the sub-sample of households in which the respondent *does* personally hold insurance. Of course, this restriction would imply that it is no longer possible to draw any inferences regarding insured versus uninsured households. And there would still be under-reporting of a specific kind: Households with both types of insurance but in which the respondent only has one type, will be confounded with households with only one type of insurance. But arguably, measurement error is less of a concern with this restriction. In effect, the choice between funeral and life insurance is now observed at the household level.

In Table 2.5, columns 1–3 present the results of regressing life insurance take-up on the same variables and controls as before, but where the sample is restricted to households in which the respondent has personally taken out insurance. The results clearly indicate that given insurance take-up, increasing asset ownership conditional on income is associated with life insurance rather than funeral insurance. Columns 4–6 repeat the exercise for holders of formal insurance only, i.e. it compares life insurance to formal-sector funeral insurance. The qualitative interpretation does not change, though in column 4 the coefficient on home ownership is not significant. Columns 7–9 focus on the black and coloured subpopulation, with the same qualitative result.

In summary, the results of the empirical analysis are supportive of one of the predictions from theory. For all the specifications, it was found that asset ownership is positively related to life insurance take-up, as predicted. The negative relation between asset ownership and funeral insurance was not always significant, but neither

was it found to be positive. However, data limitations has not allowed some of the key predictions of the model to be tested.

2.7 Discussion and conclusion

The main contribution of this paper is the attempt to make precise the role of funeral insurance as distinct from life insurance. A model is laid out which abstracts from the organisational form of the insurance provider and focuses on a key difference between the two insurance types: The payout from a funeral insurance can be spent only on the funeral, whereas the payout from life insurance can be disposed of as the beneficiary pleases. The parent puts a higher weight on funeral expenditure than does the child. The parent may use funeral insurance to impose a minimum level of expenditure on her own funeral, but this is costly since insurance premia are actuarially unfair. Life insurance is identically priced but is preferred at higher wealth levels because less of it is required to obtain a certain level of expected utility. At very high levels of wealth, there is no demand for life or funeral insurance.

The central claim made is that the raison d'être of funeral insurance is to solve an inter-generational conflict over resource allocation. In relation to the literature on life insurance and bequests, the model highlights the possibility that even if bequests are voluntary they need not be altruistic in the usual sense of the word.

Today's South Africa is perhaps unique in that funeral insurance (formal and informal) and life insurance are both widespread. It is therefore the ideal setting for testing the theory of this paper. The empirical findings are broadly consistent with one of the model's predictions, but data limitations precludes testing of some of its key implications.

Within the broad category of literature on informal finance, this paper has close parallels to the work on ROSCAs by Besley, Coate & Loury (1993). They too assume a world without credit. Like in this paper, a typically group-based financial contract is contrasted with a standard financial product. In their case, ROSCAs are shown to be ex-ante preferable to pure savings schemes for financing durables. Another parallel is that ROSCAs are typically associated with the informal sector, but that informality is hardly the defining attribute. A thorough understanding of ROSCAs is gained only by studying under what circumstances the particular contract is preferable to individual savings accounts, informal or not.

Case et al. (2008) defend the view that "social norms are held strongly and play an important role in setting funeral spending", based both on their reading of the ethnographic literature and on their own experience in training local field workers. Roth (2001), in addition to belief in afterlife, also points to social mechanisms when explaining high funeral costs: elaborate funerals impart gravity to the meeting of relatives which takes place afterwards, they demonstrate family dignity and double as important social events in which "conspicuous consumption [is] clearly in evidence".

However, the notion that the immediate family of the deceased faces social pressure to spend substantial resources on the funeral does not contradict the view presented in this paper. The central claims made here are (1) that there may be a conflict of interest over how much to spend on funerals, with the deceased's preferred level of expenditure exceeding that of his or her heirs, and (2) that funeral insurance can serve to constrain the survivors' choices in this regard. But there are reasons to believe that the wider family or social group may also prefer a higher level of funeral expenditure than does the deceased's immediate family. It is plausible that extended family and social relations benefit more from being related to or associated with someone of a high social status and a dignified funeral, than they do from the private consumption of the heirs of the deceased. Of course, they also benefit directly from the food and hospitality provided at the funeral.

So the deceased and her friends and extended family may share a preference for a larger funeral expenditure than the immediate family would prefer. Whereas friends and extended family are alive and can apply social pressure on the heirs, the deceased person may achieve a similar effect by purchasing funeral insurance while alive.

Funeral insurance has been documented through much of recorded history and is still popular in many parts of Africa. Though funeral societies have long been known to historians and has more recently started to receive the attention of economists and other social scientists, there has been little emphasis on the particularities of the insurance product they offer. Given its historical and contemporary importance, this is striking. It may appear that the process of economic development entails a transition from a world of informal or semi-formal funeral societies to one in which formal life insurance predominates. As such it is proper that substantial effort has gone into understanding contractual limitations of informal risk-pooling groups and their consequences for household behaviour. But arguably, our understanding of risk coping in developing countries would remain incomplete if we do not also understand the dynamics of insurance product demand.

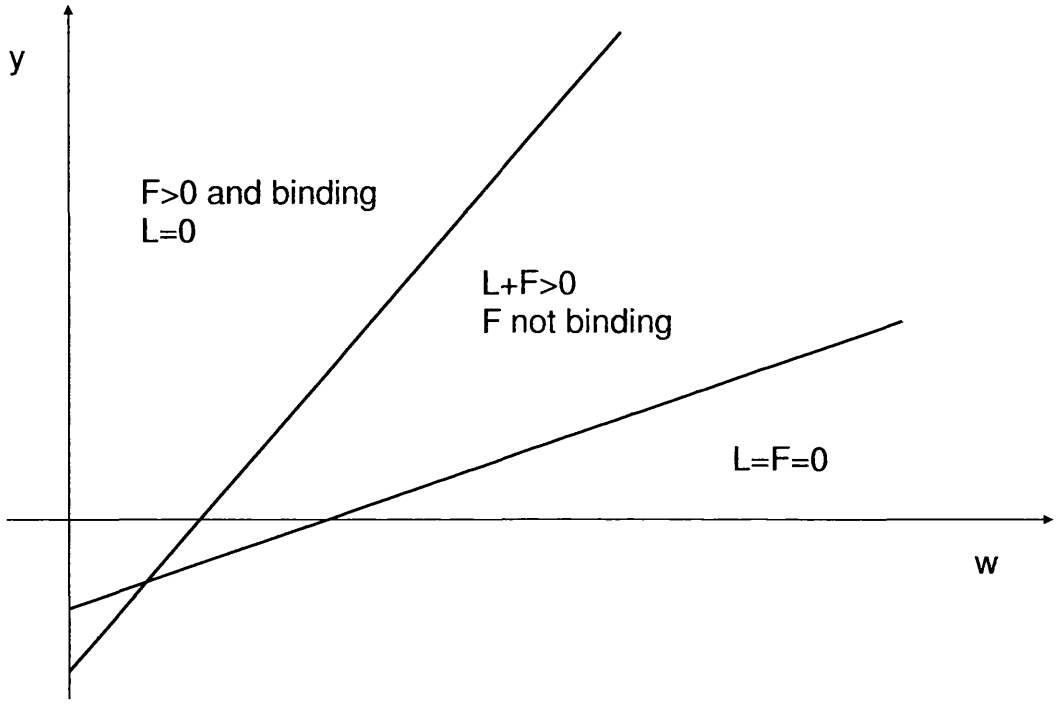

Figure 2.1: Illustration of proposition

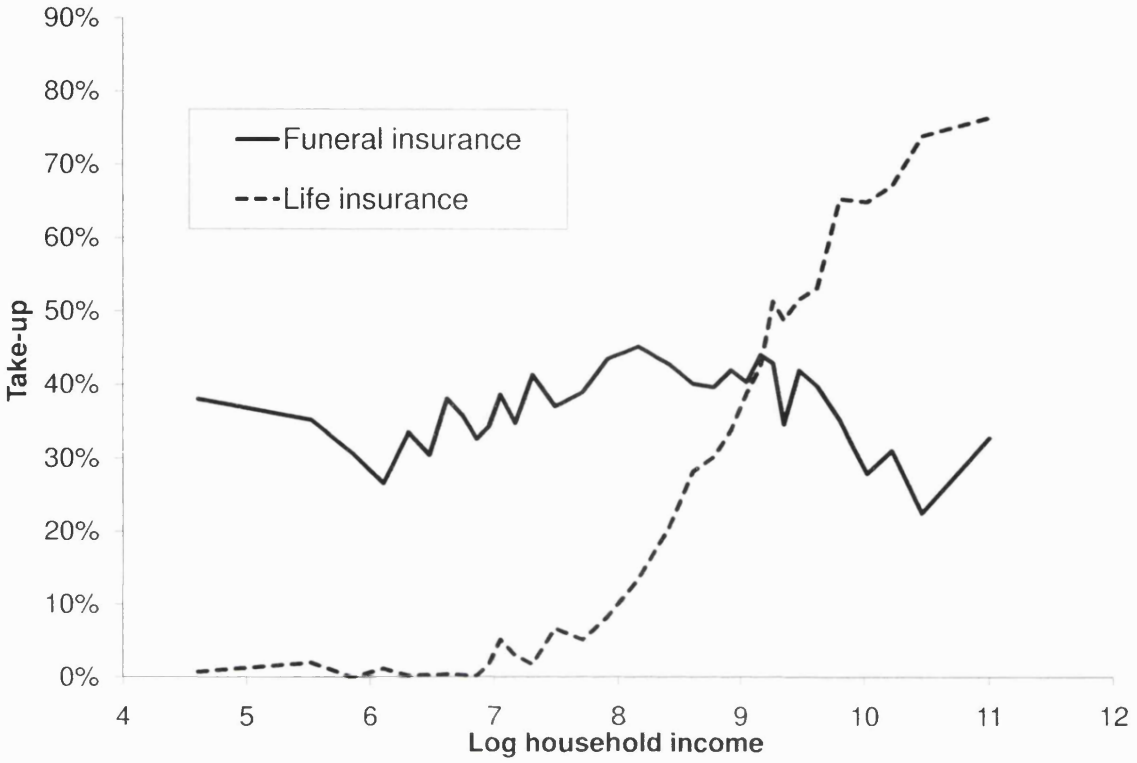

Figure 2.2: Take up of life and funeral insurance amongst respondents with own children

Table 2.1: Summary statistics (weighted).

Variable	Mean	Standard deviation
Respondent funeral society membership	0.22	0.41
Respondent has formal funeral insurance	0.12	0.32
Respondent has life insurance	0.11	0.32
Respondent age	37.1	16.1
Monthly household income (ZAR)	4146	6617
Household has car	0.27	0.45
Household owns 'proper' house	0.19	0.39
Rural	0.41	0.49
Observations	24489	

Table 2.2: Regression results, main specification

	Regression 1		Regression 2		Regression 3	
Dependent variable	**Funeral insurance**	**Life insurance**	**Funeral insurance**	**Life insurance**	**Funeral insurance**	**Life insurance**
Household has car	-0.0751	.459***	-.157***	.43***	-.141**	.409***
	[.058]	*[.042]*	*[.0597]*	*[.0425]*	*[.0603]*	*[.0425]*
Household owns 'proper' house	-0.0456	.0897*	-.114*	0.0523	-.0949*	0.0583
	[.0457]	*[.053]*	*[.0587]*	*[.045]*	*[.0533]*	*[.0438]*
Log household income	0.0525	.598***	0.0403	.595***	0.0371	.585***
	[.0391]	*[.0348]*	*[.0349]*	*[.0257]*	*[.0367]*	*[.0261]*
Age, rural and province controls	No		Yes		Yes	
Controls for own children and married/cohabiting status	No		No		Yes	
Observations	24489		24489		24489	

Weighted bivariate probit regression. The dependent variables are binary take-up indicators for funeral and life insurance. The respondent is coded as having funeral insurance if he/she has funeral insurance with a formal provider or if he/she attends a funeral society meeting weekly or monthly. The age controls are binary variables indicating that the respondent is aged 20 or over, 30 or over, etc. up to 80 or over. Robust standard errors, shown in brackets, are clustered at the province level. *** significant at 1%, ** significant at 5%, * significant at 10%.

Table 2.3: Regressions results, blacks and coloureds only

	Regression 1		Regression 2		Regression 3	
Dependent variable	**Funeral insurance**	**Life insurance**	**Funeral insurance**	**Life insurance**	**Funeral insurance**	**Life insurance**
Household has car	.0996*	.294***	0.0694	.285***	0.0649	.265***
	[.0523]	*[.0538]*	*[.0513]*	*[.0463]*	*[.0519]*	*[.0426]*
Household owns 'proper' house	0.036	.178*	-0.0659	0.101	-0.0554	0.109
	[.0505]	*[.0933]*	*[.0588]*	*[.0846]*	*[.0592]*	*[.0833]*
Log household income	.111***	.68***	.104***	.672***	.0986***	.666***
	[.0423]	*[.0403]*	*[.0347]*	*[.0375]*	*[.0375]*	*[.0373]*
Age, rural and province controls	No		Yes		Yes	
Controls for own children and married/cohabiting status	No		No		Yes	
Observations	24489		24489		24489	

Weighted bivariate probit regression. The dependent variables are binary take-up indicators for funeral and life insurance. The respondent is coded as having funeral insurance if he/she has funeral insurance with a formal provider or if he/she attends a funeral society meeting weekly or monthly. The age controls are binary variables indicating that the respondent is aged 20 or over, 30 or over, etc. up to 80 or over. Robust standard errors, shown in brackets, are clustered at the province level. *** significant at 1%, ** significant at 5%, * significant at 10%.

Table 2.4: Regressions results, formal insurance only

	Regression 1		Regression 2		Regression 3	
Dependent variable	**Funeral insurance**	**Life insurance**	**Funeral insurance**	**Life insurance**	**Funeral insurance**	**Life insurance**
Household has car	0.114	.453***	0.017	.427***	0.022	.407***
	[.0964]	*[.043]*	*[.112]*	*[.0428]*	*[.114]*	*[.042]*
Household owns 'proper' house	.117**	.101**	-0.0351	0.058	-0.0244	0.0638
	[.0569]	*[.051]*	*[.0388]*	*[.0446]*	*[.0354]*	*[.0432]*
Log household income	.288***	.601***	.283***	.595***	.279***	.586***
	[.0352]	*[.0352]*	*[.0381]*	*[.0256]*	*[.0388]*	*[.0261]*
Age, rural and province controls	No		Yes		Yes	
Controls for own children and married/cohabiting status	No		No		Yes	
Observations	24489		24489		24489	

Weighted bivariate probit regression. The dependent variables are binary take-up indicators for formal funeral and life insurance. The respondent is coded as having funeral insurance if he/she has funeral insurance with a formal provider or if he/she attends a funeral society meeting weekly or monthly. The age controls are binary variables indicating that the respondent is aged 20 or over, 30 or over, etc. up to 80 or over. Robust standard errors, shown in brackets, are clustered at the province level. *** significant at 1%, ** significant at 5%, * significant at 10%.

Table 2.5: Regressions results, insurance holders only

	1	2	3	4	5	6	7	8	9
Dependent variable	**Life insurance**	**Life insurance**	**Life insurance**	**Life insurance**	**Life insurance**	**Life insurance**	**Life insurance**	**Life insurance**	**Life insurance**
Household has car	.597***	.598***	.559***	.496***	.541***	.513***	.317***	.31***	.288***
	[.0708]	*[.0727]*	*[.0752]*	*[.13]*	*[.126]*	*[.125]*	*[.0798]*	*[.0783]*	*[.0758]*
Household owns 'proper' house	.22***	.189***	.182***	0.0193	.114**	.11**	.269***	.214***	.216***
	[.0684]	*[.0405]*	*[.0427]*	*[.0811]*	*[.0546]*	*[.0557]*	*[.084]*	*[.0624]*	*[.063]*
Log household income	.752***	.72***	.712***	.559***	.509***	.503***	.76***	.72***	.714***
	[.0452]	*[.0292]*	*[.0306]*	*[.0486]*	*[.0579]*	*[.0584]*	*[.0551]*	*[.0427]*	*[.0431]*
Sub-sample	All insurance holders	All insurance holders	All insurance holders	Formal insurance holders	Formal insurance holders	Formal insurance holders	Black/ coloured insurance holders	Black/ coloured insurance holders	Black/ coloured insurance holders
Age, rural and province controls	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Controls for own children and married/cohabiting status	No	No	Yes	No	No	Yes	No	No	Yes
Observations	9712	9712	9712	6689	6689	6689	5793	5793	5793

Weighted probit regression restricted to the subsample of insurance holders. The dependent variable is the binary take-up indicator for life insurance. The age controls are binary variables indicating that the respondent is aged 20 or over, 30 or over, etc. up to 80 or over. Robust standard errors, shown in brackets, are clustered at the province level. *** significant at 1%, ** significant at 5%, * significant at 10%.

A Proof of the proposition

Solve backwards from the child's problem. It is straightforward to show that as long as the funeral insurance payout does not impose a binding constraint on funeral expenditure,

$$\begin{aligned} f &= \gamma(W-a) \\ c &= (1-\gamma)W+\gamma a. \end{aligned}$$

On the other hand, if the payout *is* binding, then clearly

$$\begin{aligned} f &= F \\ c &= W-F. \end{aligned}$$

Since funeral insurance only imposes a *lower* bound on funeral expenditure, and since the logarithmic utility function satisfies the usual conditions, the solution must equal the child's unconstrained preference if this exceeds funeral insurance payout and exactly equal the funeral insurance payout otherwise:

$$\begin{aligned} f(W,F) &= \max(\gamma(W-a),F) \\ c(W,F) &= W-f(W,F) \end{aligned}$$

Moving on to the parent's problem, note that as long as funeral insurance is not binding, it is equivalent to life insurance. Furthermore when funeral insurance *is* binding, there will be no life insurance. This is because as long as funeral insurance payout is binding on the child, any life insurance payout would be spent entirely on the child's personal consumption from which the parent derives no utility. Thus it is sufficient to focus on three cases: No insurance, life insurance alone and binding funeral insurance alone.

Without any insurance, the parent's expected utility is simply:

$$\begin{aligned} V_N &= q\ln(C)+(1-q)\ln(f(W,F)) \\ &= q\ln(w_s)+(1-q)\ln(\gamma(w_d-a)) \\ &= q\ln(w+y)+(1-q)\ln(\gamma(w-a)) \end{aligned}$$

Assume next that the parent takes out only life insurance. Then her expected utility is

$$V_L = q\ln(w_s)+(1-q)\ln(\gamma(w_d-a))$$

with

$$\begin{aligned} w_s &= w + y - \alpha(1-q)L \\ w_d &= w + (1 - \alpha(1-q))L. \end{aligned}$$

The first-order condition yields:

$$L = \frac{-(\alpha-1)w + (1-\alpha(1-q))y + \alpha q a}{\alpha(1-\alpha(1-q))}$$

Thus L increases in y but decreases in w. Non-negativity requires

$$y > \frac{(\alpha-1)w - \alpha q a}{1-\alpha(1-q)};$$

if this does not hold then life cover must be zero.

The second-period wealth levels corresponding to the interior solutions with positive life insurance are:

$$\begin{aligned} w_s &= q\frac{w + (1-\alpha(1-q))y - (1-q)\alpha a}{1-\alpha(1-q)} \\ w_d &= \frac{w + (1-\alpha(1-q))y + \alpha q a}{\alpha}, \end{aligned}$$

and expected utility with optimal, positive life insurance becomes:

$$\begin{aligned} V_L &= q\ln(w_s) + (1-q)\ln(\gamma(w_d - a)) \\ &= q\ln q + \ln(w + (1-\alpha(1-q))y - (1-q)\alpha a) \\ &\quad -q\ln(1-\alpha(1-q)) + (1-q)\ln\gamma - (1-q)\ln\alpha \end{aligned}$$

Next assume that the parent takes out only funeral insurance, and that it is binding. Expected utility is then

$$V_F = q\ln(w_s) + (1-q)\ln(F).$$

The first-order condition gives

$$F = \frac{w+y}{\alpha}.$$

From this one may compute the next-period wealth levels and optimal expected utility for the case of binding funeral insurance:

$$\begin{aligned} w_s &= q(w+y) \\ w_d &= \frac{1+\alpha q}{\alpha}w + \frac{1-\alpha(1-q)}{\alpha}y \\ V_F &= q\ln q - (1-q)\ln\alpha + \ln(w+y) \end{aligned}$$

The optimal F clearly always satisfies non-negativity, but it was also assumed that the insurance is binding on the child's funeral expenditure. That would require

$$F > \gamma (w_d - a) = \gamma \left(\frac{1 + \alpha q}{\alpha} w + \frac{1 - \alpha (1 - q)}{\alpha} y - a \right)$$

which yields the condition

$$y > \frac{(\gamma (1 + \alpha q) - 1) w - \gamma \alpha a}{1 - \gamma (1 - \alpha (1 - q))}.$$

The optimal expected utility functions for each of the three main cases (life insurance only, binding funeral insurance only, and no insurance) have been determined. Now the overall solution to the parent's problem given a set of parameters may be found as follows: First check whether the non-trivial conditions for interior solutions are satisfied. (Life insurance must be non-negative, and funeral insurance must be binding on the child for the above solution to hold. The no-insurance solution is always feasible.) Then compare the optimal expected utilities for the feasible cases and select the greatest.

Note, however, that the life insurance solution will always dominate the no-insurance solution as long as optimal life insurance is positive. This is because no insurance is essentially a special case of the life-only scenario with $L = 0$.

Thus there are four conditions that need to be taken into account when determining the solution for a given set of parameters. They are:

$$\begin{aligned} V_F &> V_L \\ V_F &> V_N \\ L &> 0 \\ & \quad F \text{ binding} \end{aligned}$$

Using the above results and rearranging in terms of y, these may be written

$$\begin{aligned} y &> \frac{\kappa w - a}{1 - \kappa}, \text{ with } \kappa = \frac{\gamma - (\gamma (1 - \alpha (1 - q)))^q}{\gamma \alpha (1 - q)} < 1 \\ y &> \frac{(1 - k) w - a}{k}, \text{ with } k = \frac{q^{\frac{q}{1-q}}}{\gamma \alpha} \\ y &> \frac{(\alpha - 1) w - \alpha q a}{1 - \alpha (1 - q)} \\ y &> \frac{(\gamma (1 + \alpha q) - 1) w - \gamma \alpha a}{1 - \gamma (1 - \alpha (1 - q))}, \end{aligned}$$

respectively. All four conditions divide w, y-space in two regions along a straight line. The boundary lines all have negative intercepts proportional to a, and positive slopes. Since non-negative wealth and income are assumed, it is the first quadrant which is

of interest.

Given the above results and the stated assumptions it can be shown that along the $y = 0$ axis, the line defined by $V_F = V_L$ crosses at a lower value of w than the boundary for binding F, which in turn crosses at a lower value of w than the boundary for $L = 0$. The slope of these lines decline in the same order. Hence the three lines do not cross in the first quadrant: The $V_F = V_L$ line is everywhere to the left of the boundary for binding F, which means that the binding F solution is feasible whenever its value dominates the non-binding solution.

The boundary for binding F in its turn is everywhere to the left of the $L = 0$ boundary. Therefore, F is binding whenever $V_F > V_L$, so $V_F = V_L$ is the effective boundary between a binding-F region and an L region. And since the boundary for binding F is to the left of $L = 0$, funeral insurance can never bind where optimal L is zero. This means that the parent will not take out any insurance to the right of $L = 0$.

Therefore, $V_F = V_L$ and $L = 0$ divide the first quadrant in three regions which may be described as follows. Holding y fixed at any non-negative level and increasing w from zero, there is first a region in which binding funeral insurance of increasing quantity is preferred. (Due to the negative intercepts, this region is of positive measure even for $y = 0$.) As w increases further, a boundary is crossed after which life insurance is preferred. The quantity of life insurance decreases to zero as wealth increases. Beyond this point the parent prefers not to take out any insurance.

Finally, recall that life insurance cover of value L is equivalent to any mixture of funeral and life insurance of the same total value, where the funeral cover remains below the binding level. This concludes the proof.

B Data variable construction

In the section 'Financial surveys', the survey asks: "Here is a list of different types of policies and investment plans which you can take out with a financial services company. Can you please tell me which, if any, you *PERSONALLY* have? This excludes any cover or benefits provided by your employer/company" (original emphasis). Respondents who have indicate either 'Life cover policy' or 'Endowment, investment, savings or education plan/policy *with life cover*' (original emphasis) are coded as having life insurance. Those who indicate the option 'Funeral insurance', are coded as formal funeral policy holders.

In the 'Sport, entertainment and leisure' section, the following question is asked: "For the activities listed below, please indicate your *personal frequency of each activity*, if at all" (original emphasis). The options given were 'weekly', 'monthly', 'yearly' and 'not at all'. Respondents reporting that they attend funeral society meetings 'weekly' or 'monthly' are coded as funeral society members. A respondent is coded as having funeral insurance if he/she has a formal funeral policy or is a member of a funeral

society. No information on premia, cover or conditions is provided for either type of insurance.

There are several caveats to keep in mind when interpreting these variables. First, it appears to be quite common for people to get funeral cover through their employer in South Africa, but this is explicitly excluded from the question above and therefore is not recorded in the data. Second, it is also quite common for funeral parlours to sell a form of funeral cover (illegally). These may not be recorded in the data since the survey explicitly asks for funeral cover with a "financial services company". Third, some insurance companies masquerade as funeral societies. These may not be recorded in the data since they do not require their customers to attend regular society meetings.

In the survey, monthly household income is recorded not as a number but as being in one of 33 income ranges. These are coded to a scalar income variable by using the midpoint of each band. The highest band, ZAR 40,000 and above, is coded as ZAR 60,000. 1% of households report this income band.

The respondents are asked "How many motor cars, station wagons, bakkies and mini-buses/kombis, including company cars, are there in your household? Please exclude motor bikes, scooters and trucks." The household was coded as having a car if the response was one or more.

The household is coded as owning a 'proper' home if all of the following are satisfied: It owns its dwelling. The dwelling is a 'House', 'Cluster House', 'Town House' or 'Flat', as opposed to a 'Matchbox House', 'Traditional Hut', 'Hostel', 'Hotel / Boarding House', 'Compound', 'Room in Backyard', 'Squatter Hut', 'Caravan' or 'Other'. The dwelling is electrified and has water on tap, a hot water tank, a toilet and a sink.

The 'Rural' indicator is set to 1 if the household's community size is reported to be "Less than 500/Rural".

The respondent is coded as having own children if he/she answered yes to the question: "Do you have any young or unmarried children of your own?" There is no information on children more generally.

The respondent is coded as married/cohabiting if he/she reported "Married or living together" under marital status.

3 Liquidity constraints: A robust analysis using panel data

3.1 Introduction

Credit constraints are a standard assumption in development economics, but there is surprisingly little robust evidence that such constraints are present and binding on household behaviour. Furthermore, many observations that are consistent with liquidity constraints are equally consistent with precautionary saving or a simple Keynesian model where consumption varies with current income. Recent work on credit constraints in rich countries has exploited the idea that under standard assumptions and perfect capital markets, there should be no jump in consumption associated with an anticipated increase in income. This paper applies that methodology to a large panel data set of black South African households. Having established that expenditure jumps are present, further evidence is presented indicating that the jumps are more likely to be due to liquidity constraints than to precautionary saving behaviour or Keynesian consumption.

Credit market imperfections are central to current development economics theory. And in models of household behaviour in developing countries, such imperfections are frequently operationalised as a constraint on the amount of credit available to the household (often taken to be zero). But given how often credit constraints are assumed, there is little evidence of their presence. In his entry on development economics in the New Palgrave Dictionary of Economics, Ray (2008) discusses credit market failure and constrained credit as a key theme of the field, but writes that "[t]he direct empirical evidence on the existence of credit constraints is surprisingly sparse". Banerjee (2003) surveys the evidence for capital market imperfections in developing countries. In this context, perhaps the most telling piece of evidence he presents is the often wide dispersal between interest rates on saving and borrowing in many parts of the developing world. But as Banerjee & Duflo (2005) write, this is only suggestive of credit constraints.[17]

In fact, many observations that are consistent with credit constraints are explained equally well by Keynesian consumption or precautionary saving behaviour. In particular, this applies to observations that consumption responds to anticipated increases in income, as well as to observations of 'excess sensitivity' of consumption with respect to temporary income shocks which have only a small impact on permanent income. Apart from the direct interest in establishing whether credit constraints are present, this relates to the long-standing debate on whether poor household optimise within their constraints ("poor but rational") or not (Duflo 2006). In particular, finding

[17] Morduch (1990) tests for liquidity constraints in a developing country context in his unpublished PhD thesis.

that the consumption behaviour of poor households is not guided by forward planning could have large consequences for the study of development.

This paper uses data on a large panel of black South African households to investigate the response of household expenditure to a large, anticipated increase in income. The public, non-contributory old age pension scheme is widely known in the population, and take-up amongst Africans is high and largely predicted by age. An approximate Euler equation is derived, from which it is clear that a jump in expenditure associated with an anticipated increase in income is as consistent with precautionary saving as with a binding liquidity constraint. Furthermore, the consumption jump can also be explained by Keynesian consumption. The prediction that consumption should not respond to the anticipated income change is rejected, even when allowing for household-specific trends. Furthermore, evidence is presented that household savings increase when a member crosses the age-qualifying age, making the precautionary saving story less likely. Finally, a change in the age at which the child support grant lapses is used to robustly identify an anticipated decrease in household income. This allows for a test of liquidity constraints versus Keynesian consumption behaviour. Liquidity constraints cannot be rejected in favour of Keynesian consumption, though the power of the test may be problematic.

The main contributions of this paper are as follows. It is one of a small number of studies that provide explicit tests of the standard consumption model with perfect capital markets using expenditure data from a developing country. Uniquely, it relies on large panel data set and a large and highly anticipated source of income to do so. The use of fixed-effects estimation and clean identification provides for a robust rejection of the standard model with perfect capital markets amongst black South Africans. This is also the first paper to attempt the distinction between liquidity constraints, precautionary saving and Keynesian consumption as explanations of expenditure jumps in a developing country. By showing that liquidity constraints is a more likely explanation of household behaviour than Keynesian consumption, the paper also contributes to the debate on whether people in developing countries are 'poor but rational'.

The remainder of the paper is organised as follows. The next section provides an overview of related literature. In the theory section, an approximate Euler equation which incorporates liquidity constraints and a precautionary saving motive is derived from a standard consumption model, and the relevant implications for empirical work are drawn out. The evidence section provides some background on the old age pension and child support grants in South Africa and describes the data before presenting the empirical specifications and results. Finally, there is a brief discussion and conclusion.

3.2 Related literature

Hall (1978) extended the standard additive consumption model[18] to the case with stochastic earnings. The paper's influential central finding is that, conditional on today's consumption, no other information available today is informative about consumption tomorrow.[19] In other words, consumption should respond to income changes only to the extent that these reflect changes in permanent income. This prediction has subsequently been exploited by many economists to construct tests of the perfect capital markets assumption. One line of research looks at the effect on consumption of anticipated or unanticipated changes in income (see Browning & Lusardi (1996) for a review). Recently, two papers have argued that households in the United States and Spain adjust their consumption in response to significant and regular changes in annual income, but not in response to irregular and small payments because the computational cost of doing so outweighs the potential utility gain. Hsieh (2003) exploits the annual payout of petroleum dividends to residents of Alaska as an anticipated change in permanent income. He does not find excess sensitivity of consumption related to the petroleum payout, but on the other hand he confirms earlier results when he finds that the same households do over-react to tax refunds, which are smaller and harder to predict on average. Similarly, Browning & Collado (2001) find that the consumption patterns of Spanish households working in sectors with regular bonus payments do not differ significantly from those of households in other sectors. The findings in this paper go the other way, indicating that at least in one developing country context, the standard model with perfect capital markets is rejected even with respect to very large and highly predictable increases income.

Altonji & Siow (1987) were the first to exploit the prediction that forward-looking but liquidity-constrained agents should react differently to positive and negative changes in income. This paper follows theirs in referring to models where consumption varies with current income as 'Keynesian', though elsewhere in the literature this is also referred to myopia. Flavin (1985) tests liquidity constraints versus Keynesian consumption on US data, but relies on an instrumental variable technique and aggregate data rather than a source of anticipated income in a panel of households. Deaton (1992) investigates how closely simple 'rules of thumb' based can approximate the optimal solution to the standard consumption model. These rules of thumb typically depend positively on current income, and are hence a form of Keynesian consumption as defined here.

Deaton (1991) studies savings behaviour in the presence of liquidity constraints, and uses simulations to study the effectiveness of precautionary saving in smoothing

[18] These models assume intertemporally additive utility, constant discount factors and consumers who maximise (expected) present discounted utility over the remainder of the life-time. I follow Browning & Lusardi (1996) in referring to these as standard consumption models, though their origins lie in Modigliani's life-cycle and Friedman's permanent income hypothesis models.

[19] Parker (1997) presents consumption model which combines Hall (1978)'s framework with the durable-goods model of Mankiw (1982), and derives similar predictions.

household consumption.

One of the few articles to provide a direct test of a prediction from the standard consumption model in a developing country is Paxson (1992). The starting point for her influential paper is the prediction that the propensity to consume out of unexpected income should be much smaller than that for expected income. In other words, a large proportion of windfall gains should be saved. In line with the prediction, she finds that saving amongst agricultural households in Thailand is responsive to positive income shocks resulting from rainfall variation. Rosenzweig & Wolpin (1993) find evidence of credit constraints among farmers in India, and Banerjee & Duflo (2008) do the same for Indian firms.

Development economists have given much more attention to how limitations in contract enforcement prevents full inter-household risk-pooling between forward-looking but liquidity-constrained poor consumers. Townsend (1994) inspired a set of papers testing for and mostly rejecting full risk-sharing across households. These papers effectively assume imperfect capital markets, since otherwise there would be no need to smooth across households. Ligon, Thomas & Worrall (2002) is a recent contribution to this tradition.

One strand of the literature uses direct survey questions to elicit credit constraints. Typically, respondents are flagged as credit constrained if they report having tried to borrow but were denied credit, or if they report being discouraged from applying for credit because they felt that they would be rejected. For instance, Barham, Boucher & Carter (1996) uses this type of survey information to look at whether the presence of credit co-operatives in regions of Guatemala relaxes credit constraints. Rasmussen (2002) employs a similar methodology to data from South Africa. She finds that a high proportion of the households in her sample are credit constrained.

This is not the first paper to exploit the South African old age pension as a source of exogenous income variation. Case & Deaton (1998) look at how various expenditure categories respond to the pension income, but they do not analyse total expenditure, nor do they focus on liquidity constraints. Edmonds (2006) finds that receiving the pension leads to a decrease in child labour and an increase in school attendance and interprets this as evidence of liquidity constraints, but does not consider Keynesian consumption as an alternative explanation. Duflo (2003) finds that young girls living with a pension-eligible grandmother are in better health. The latter two both argue that the their findings are indicative of credit constraints. But their use of cross-section data sets makes it difficult to rule out certain other possible explanations, including a re-allocation of family resources when the income increases, or indeed Keynesian consumption, as alternatives to constrained liquidity.

3.3 Theory

3.3.1 A consumption model with liquidity constraints

The following derivation of an Euler equation from an additive consumption model with liquidity constraints follows Browning & Lusardi (1996). Though it will become clear that the resulting equation is approximate and that not all predictions are unambiguously related to the underlying model parameters, it has two attractive properties. First, it is (approximately) linear when a CRRA utility function is specified. Second, it is fairly general in that it relates the consumption pattern to household characteristics, impatience, the interest rate and general uncertainty (allowing for precautionary saving) as well as to liquidity constraints.

An agent lives for T periods. In each period t she sets current consumption c_t so as to maximise

$$E_t \sum_{j=0}^{T-t} \delta^j u\left(c_{t+j}, Z_{t+j}\right)$$

subject to

$$\begin{aligned} w_\tau &= \left(1 + r_\tau\right) w_{\tau-1} + y_\tau - c_\tau \\ w_\tau &\geq 0 \end{aligned}$$

for $t \leq \tau \leq T$. Here, $\delta \in (0,1)$ is a discount factor, and the instantaneous utility function u satisfies $u' > 0$ and $u'' < 0$. Utility is a function both of consumption and of other characteristics Z_t. The first constraint defines the evolution of wealth: Wealth at the end of period τ, w_τ, is the wealth remaining at the end of the previous period $w_{\tau-1}$, augmented by the interest rate r_τ, and adjusted for current-period labour income y_τ and consumption. The second constraint is the liquidity constraint: Non-negative wealth at the end of every period implies that the agent cannot borrow. However, note that the following analysis is robust to allowing some borrowing, i.e. a negative rather than zero lower bound on wealth.

Uncertainty may be present in future labour income, household characteristics or interest rate, but the model is agnostic about its form.

The first-order conditions with respect to consumption yield the Euler equation

$$u'\left(c_t, Z_t\right) = E_t \delta \left(1 + r_{t+1}\right) u'\left(c_{t+1}, Z_{t+1}\right) + \lambda_t,$$

where u' denotes the derivative of the instant utility function with respect to the first argument (consumption). The Kuhn-Tucker multiplier λ_t is always non-negative, and positive when the liquidity constraint binds in period t. An immediate consequence of this equation is that ceteris paribus, a binding liquidity constraint tends to reduce current consumption.

In order to obtain more specific predictions, assume a CRRA utility function

$$u(c_t, Z_t) = \frac{\left(\frac{c_t}{e^{aZ_t}}\right)^{1-\rho}}{1-\rho}$$

for constants ρ and a. Note that the role of Z in this specification is to scale consumption, like in a conversion to adult equivalents.

With this specification, the Euler equations become

$$c_t^{-\rho}\left(\frac{1}{e^{aZ_t}}\right)^{1-\rho} = E_t\delta\left(1+r_{t+1}\right)c_{t+1}^{-\rho}\left(\frac{1}{e^{aZ_{t+1}}}\right)^{1-\rho} + \lambda_t,$$

which may be written

$$\delta\left(1+r_{t+1}\right)\left(\frac{c_{t+1}}{c_t}\right)^{-\rho}\left(e^{a(Z_{t+1}-Z_t)}\right)^{\rho-1} = 1 - \lambda_t c_t^{\rho}\left(e^{aZ_t}\right)^{1-\rho} + \nu_{t+1}$$

with $E\nu_{t+1} = 0$. Even if the uncertainty in the model was fully specified, it would be difficult to derive the precise distribution of ν_t. To do so would require a solution for c_t, which is not in general available analytically. In any case, the variance of the error term ν_{t+1}, which will be denoted σ_{t+1}^2, is a function of the underlying uncertainties and is a measure of how these uncertainties map into the unpredictability of future consumption.

By rearranging, reindexing by subtracting 1 from all time indices, taking logs on both sides and using Taylor expansion approximations to logarithms, one obtains

$$\ln c_t - \ln c_{t-1} \approx \frac{\ln\delta}{\rho} + \frac{r_t}{\rho} + \frac{a(\rho-1)}{\rho}\left(Z_t - Z_{t-1}\right) + \frac{\sigma_t^2}{2} + \lambda_{t-1}c_{t-1}^{\rho}\left(e^{aZ_{t-1}}\right)^{1-\rho} + \frac{\lambda_{t-1}^2 c_{t-1}^{2\rho}\left(e^{aZ_{t-1}}\right)^{2(1-\rho)}}{2}. \tag{3.1}$$

where ε_t is a higher-order remainder term with $E\varepsilon_t = 0$.

This approximate equation relates changes in (log) consumption to various parameters of the model. The first term on the right-hand side shows that impatience (lower δ) has the expected effect of increasing consumption today relative to that of tomorrow. Interest has, again as expected, the opposite effect. The third term on the right-hand side captures the general notion that changes in other characteristics, such as household demographics, may explain changes in the household's consumption profile over time. The fourth term on the right-hand side predicts precautionary consumption behaviour: The more uncertainty there is, the greater the propensity to save for a rainy day, and this depresses current consumption relative to future consumption. (Consumption increases over time as the agent's buffer stock of savings builds up in expectation.) The next two terms are increasing in λ_{t-1} and show that a binding liquidity constraint to depress consumption in that period and hence increase consumption growth.

3.3.2 Implications

Several previous empirical tests of the standard consumption model with perfect credit markets rely on the prediction that an anticipated change in income should not in itself have any effect on consumption. Rejections of this prediction have been taken as evidence of liquidity constraints. In the light of the Euler equation above, which explicitly incorporates the possibility of liquidity constraints, it is clear that these constraints are indeed a possible explanation for a change in consumption levels associated with a change in income. A household that expects an increase in income would, in the absence of liquidity constraints, like to smooth its consumption by borrowing before the income increase happens. But a liquidity constraint may prevent the increase from happening until the higher income is realised. In terms of the equation, recall that a binding liquidity constraint is associated with $\lambda_t > 0$, and that $\lambda_t = 0$ corresponds to the absence of a binding constraint. Clearly a binding constraint is associated with a steeper increase in income.

The model rules out borrowing, but not saving (otherwise the household would simply consume its income in each period). This is probably realistic, though the return on savings may be low and uncertain. When a household faces an anticipated decrease in income, the relevant technology is savings and not credit. This gives rise to an asymmetry of predictions: In the presence of constraints on liquidity (but not on savings), an anticipated *increase* in income may be associated with a contemporaneous increase in consumption. But an anticipated *decrease* income should not be accompanied by a correspondingly steep drop in consumption, because the household is able to smooth its consumption profile using savings.

More generally, a change in the level of consumption in a household can have several possible causes. These causes may be loosely classified as 'within-trend', 'trend-breaking' or 'model-breaking'. Even without any changes in the right-hand side variables, consumption may be on an increasing or decreasing trend due to the effect of impatience (which tends to put the household on a decreasing consumption trend), interest (saving to achieve higher consumption tomorrow) or uncertainty (savings as a buffer).

Consumption may also be breaking out of its trend, in response to changes in the right-hand side variables. The possible causes are a change in the rate of return on savings, changes in 'other relevant characteristics' such as family demographics or a change in the level of uncertainty about future consumption.

Liquidity constraints can come under either of these two categories: In the presence of liquidity constraints, current consumption tends to be suppressed, implying an increasing trend in consumption. But liquidity constraints could also move from being binding to non-binding and vice versa, which would have a 'trend-breaking' effect on consumption.

The third and final category of explanations for changes in the consumption level

is that the model is misspecified. It could be that the 'true' utility function is not CRRA, or that Z enters the model in a more complex manner than the above allows for. Or it could be a rejection of the (arguably) core assumptions of the model, namely that utility is temporally additive and that consumers are expected-utility maximisers. Keynesian consumption violates even these most basic assumptions of the standard model. Extreme Keynesian consumers do not plan ahead at all. Instead, they simply consume a fixed fraction of current income. Considered as a theory of household behaviour, extreme Keynesian consumption is simplistic. But it does capture a sense of myopia or lack of forward planning which can serve as a benchmark against which to measure the meticulously planning households of the standard models. Specifically the consumption of a Keynesian household will track that of its income, regardless of whether income increases or decreases, and regardless of whether the change is anticipated or not.

The aim of the rest of the paper is to use an anticipated source of income to examine household data for changes in consumption behaviour, and attempt to distinguish between the possible theoretical explanations for these.

3.4 Evidence

3.4.1 The old age pension and the child support grant

The South African old age pension is a non-contributory, means-tested public benefit scheme, originally set up in the 1920s to provide a retirement income for the minority of white workers who did not have pension arrangements through their employer. In 1989, the government committed to removing racial inequalities in the system, and from 1993 onwards the new scheme was fully operational throughout the country. In rural areas the money is distributed from a sophisticated network of mobile paypoints reaching each beneficiary once a month. Today the pension is a major source of income amongst the poor all over the country, and in many households the main breadwinner is a pensioner. See Case & Deaton (1998) for more detail.

To qualify for the pension, the applicant has to be old enough as well as pass a means test. Only people whose wealth and income are below certain levels qualify for the grant. The means test is set at a level where most *white* elderly fail to qualify, but a large majority of *black* elderly easily pass the hurdle. Hence for black people the pension can largely be regarded as a universally available and secure monthly income stream, beginning at the lower qualifying age (60 for women, 65 for men) and continuing until death. It is independent of the economic situation of family members, except the spouse for married pensioners. From an econometric perspective, it provides a convenient source of nearly exogenous income variation amongst people around the qualifying age.

Given the popularity, reach and financial importance of the pension, it is a highly anticipated source of income. The qualifying age limits are also widely known. (Duflo

2003, Case & Deaton 1998)

The child support grant was introduced in 1998 to replace the earlier state maintenance grant. It is paid to the primary care-giver (most often the mother) of young children. It is in principle means-tested, and a large proportion of black mothers with age-eligible children qualify. From the beginning it covered children up to and including the age of six, but from 2002 there was a gradual expansion to older children. For the 2001 and 2002 surveys used in this study, the upper age limit was six. When the 2003 survey was conducted, the upper age limit was eight years. See Rosa & Mpokotho (2004) for more details on the child support grant.

3.4.2 Data

Since the year 2000, Statistics South Africa has conducted a Labour Force Survey twice a year. It is designed as a rotating panel survey, but the longest series of longitudinal household observations is substantially shorter than it could have been, due to several 'fresh starts' which involved drawing all-new samples. This study uses the September 2001, September 2002 and September 2003 waves of the survey. The surveys are nationally representative. Only households headed by a black ('African', in South African terminology) person are considered here. After dropping 175 households due to missing data on the gender, age or population group of one or more household members, there remains 36,208 households in the panel, of which 11,962 are observed twice and 6,668 are observed three times.

The expenditure variable is constructed as follows. In each of the surveys, the households were asked: "What was the total household expenditure in the last month? Include everything that the household and its members spent money on, including food, clothing, transport, rent and rates, alcohol and tobacco, school fees, entertainment and any other expenses." The answer options were: R 0–399, R 400–799, R 800–1,199, R 1,200–1,799, R 1,800–2,499, R 2,500–4,999, R 5,000–9,999, R 10,000 or more, "Don't know" or "Refuse". The monthly expenditure variable used in this paper was coded as follows: For each of the intervals, the midpoint is used. For the "R 10,000 or more" option, expenditure is set to R 15,000. For the last two response options, monthly expenditure is coded as missing. It is the logarithm of the resulting number that is used here as the measure of household expenditure.

A household is coded as age-eligible for the old age pension if there is at least one woman aged 60+, or one man aged 65+, in the household, and age-eligible for the child support grant if at least one child below the upper age limit is present in the household at the time of observation.

Descriptive statistics are presented in Table 3.1. In Figure 3.1, take-up of the old pension is plotted against the age of the oldest male and female household members. Age-eligibility appears to be a strong predictor of pension take-up; though perhaps more cleanly for women than for men.

3.4.3 Identification and main results

The baseline empirical specification is given by

$$Y_{it} = \alpha_i + \mu_t + \gamma L_{it} + X_{it}\beta + \varepsilon_{it}.$$

Here, Y_{it} is the expenditure of household i at time (survey) t. On the right-hand side, α_i is a household fixed effect, μ_t is a time (survey) dummy, and L_{it} is an indicator for whether household i receives, or is age-eligible for, the old age pension in year t. X_{it} is a vector of control variables, and ε_{it} is the error term.

The coefficient of interest is γ. Given the specification, it is only identified within households that move from a non-pension to a pension state or vice versa. Disregarding deaths and other changes in household composition (which will be controlled for through X), this will only happen when a female household member turns 60, or a male household member turns 65, in the second or third wave of the survey. In the eligibility specification the identifying assumption is that, conditional on household demographics and other controls, there is nothing special about turning 60/65 other than that one becomes eligible for the pension. This assumes that all other changes to do with 'entering retirement' are either a direct consequence of qualifying for the pension, or are controlled for in the regression.

Throughout the analysis, age-eligibility as well as actual take-up of the old age pension will be used to identify the impact of the pension income on consumption. This is due to possible endogeneity in the take-up decision: It may be that a person's decision of when to apply for the pension is co-determined with other unobserved characteristics which may also impinge on expenditure.

For this reason, it is useful to confirm that the eligibility criterion actually predicts pension take-up before moving on to the main results. Figure 3.1 has already provided a strong indication that this is the case. Table 3.2 regresses a binary variable indicating pension receipt on the eligibility dummy. In the first column, household and year fixed effects are included but no other control variables. The regression in the second column includes household demographics (gender-specific count variables for the number of family members in the age groups 0–4, 5–14, 15–24, 25–34, 35–44, 45–54 and 55+, and household size). In both these regressions, the coefficient is large (.48 and .39) and highly significant. In the third column, further count variables for the presence of older household members are included. There are variables for people aged 55+, 60+, 65+ and 70+, separately for women and men. The coefficients on female and male age-eligibility, i.e. on the variables 60+ and 65+, respectively, are large and significant (.41 and .22). Though some of the other variables accounting for older household members are also positive and significant, these two dominate. Note also that the coefficient on female age-qualification is twice as large as the coefficient on male age-qualification. It seems reasonable to conclude that age eligibility is a strong predictor of pension take-up, but that there is a non-degenerate distribution in the

age at pension take-up.

The baseline results, looking at the response in the level of household expenditure to the pension, are presented in Table 3.3. In column 1 household expenditure is regressed on a dummy variable indicating whether the household receives the old age pension (1) or not (0). There are no control variables, apart from household and survey fixed effects. The coefficient of interest is positive at .20 and highly significant. Column 2 includes controls for household demographics (the same as above). The coefficient drops to .13 but remains highly significant. In column 3, the variables controlling for the presence of older household members by age bracket are included. The coefficient drops to .11 but is still significant at the 1% level. This indicates that household expenditure responds positively to the anticipated increase in household income.

Though these results are interesting, there is a possible endogeneity problem. The decision to take up the pension, conditional on eligibility, may be correlated with household unobservables captured in the error term. The coefficient on pension receipt potentially confounds the effect of the pension income with these unobservables. The next three columns therefore presents reduced-form versions of these results, where household expenditure is regressed on age-eligibility for the pension rather than actual take-up. Given that not all age-eligible households in the sample receive the pension, the expectation is that the magnitude of the coefficients will drop somewhat, and this is indeed what happens. Column 4 replicates the regression in column 1, except that expenditure is regressed on age-eligibility rather than pension receipt. The coefficient is somewhat lower at .15 but still highly significant. Column 5 includes household controls, and though the coefficient drops to .08 the qualitative finding is unchanged. In column 6 the count variables for the presence of older household members are included. This is equivalent to splitting the household eligibility variable into female and male age-eligibility. Female age-eligibility is highly significant at .10, but male age-eligibility is no longer significant. This is consistent with the finding above that men's pension take-up is less well predicted by eligibility.

In order to interpret these results, refer back to the Euler equation (3.1). Basic household demographics are controlled for, so these cannot account for the expenditure response. Though of course it is difficult to rule out unobserved characteristics which change over time, the household fixed effects capture anything that is time-invariant at the household level. The Euler terms involving the discount factor and interest rate may be responsible for an increasing or decreasing trend in expenditure, though the year fixed effects will capture any changes that affect the sample as a whole. But if taking up the pension is associated with a decrease in consumption uncertainty (which seems likely) as well as an increase in income, then these results cannot rule out that the observed change in consumption is due to a reduced need to for precautionary saving. Finally, the finding may be explained by liquidity constraints.

Some of these concerns may be alleviated by introducing household-level trends

in consumption in the regressions. This is computationally infeasible given the size of the data set. One could allow for trends at a more aggregate level, such as survey cluster or province, but this would defeat some of the purpose of the exercise since inter-household variation in these trends are hardly implausible even at a modest level of aggregation.

A viable alternative is to first-difference the expenditure variable. Hitherto the idea has been to study changes in expenditure levels in response to the change in income. But by first-differencing the expenditure variable, it becomes possible to look at the response in expenditure *growth* instead. This would eliminate linear household consumption trends from the analysis, such as those reflecting a constant discount factor, interest rate or precautionary incentive. Referring to the Euler equation (3.1) again, differencing the equation would eliminate everything on the right-hand side which does not change over time.

On the other hand, doing so entails a loss of statistical power since the coefficient of interest can then only be identified off households which are both observed in all three surveys, and which undergo a change in eligibility status during the period of observation.

The empirical specification for the difference analysis is

$$Y_{it} - Y_{i,t-1} = \widetilde{\alpha}_i + \widetilde{\mu}_t + \widetilde{\gamma} L_{it} + X_{it}\widetilde{\beta} + \widetilde{\varepsilon}_{it},$$

where $\widetilde{\alpha}_i$ capture a household-specific trend in consumption, $\widetilde{\mu}_t$ capture economy-wide shocks to consumption growth and $\widetilde{\gamma}$, the new coefficient of main interest, identifies a change in the growth rate rather than level of expenditure, as a response to pension take-up or eligibility.

The results of this analysis are presented in Table 3.4, which is a close parallel to Table 3.3 except that the dependent variable has been first-differenced. Columns 1–3 identify a strong surge in expenditure trend associated with the take-up of the pension (the coefficients are .33, .23 and .22 and highly significant), even when household demographics and counts of older household members are included. Columns 4–6 are again reduced-form regressions where de-trended expenditure is regressed on age-eligibility rather than actual take-up. The aim is again to overcome potential endogeneity in the take-up decision, at the cost of a downward bias of the coefficients of interest. The regression in column 4 includes no controls except fixed household and year trends. The coefficient of interest is .25 and significant at the 1% level. However, when household demographics are included, the coefficient drops to .12 and is only marginally significant (i.e., at the 10% level). And when the counts of older household members are included, both female and male age-eligibility becomes insignificantly different from zero. One interpretation of these findings is that the effect found earlier is fully accounted for by household trends and changes in household demographics. But it is equally possible that there really are liquidity constraints, but that the com-

bination of first-differencing, including fixed-effects and a full set of household controls as well as relying on a reduced-form regression has reduced effective power to the point where the behavioural response to the pension income is no longer discernible.

Some readers may be concerned that a major life change ('retirement') might systematically coincide with the qualifying age for the pension and bias the results. However, for the majority of the black population who are either unemployed or employed in the informal sector, there is no well-defined concept of 'retirement' in the sense of reaching the end of an employment contract. This is illustrated by the fact that only 2.6% of black women in the 60-64 age group receive employment-related pensions, and the equivalent number for black men in the age group 65-69 is 7.2%. Withdrawal from the labour market is much more likely to *caused* by the pension income (Ranchhod 2006), or to have already happened (Bertrand, Mullainathan & Miller 2003), than to coincide with it in a way which might bias the identified effect on consumption.

3.4.4 Using household savings to test between liquidity constraints and precautionary motives

Within the general consumption theory outlined in this paper, the results so far are compatible with either liquidity constraints or precautionary saving, or both. If a household is fairly certain of a future income stream, it would be expected to smooth consumption by borrowing. On the other hand, the old age pension represents not just a sizeable income, it is also a secure income for life. Employment income is uncertain by comparison, especially in the informal labour market. Therefore, becoming eligible for the pension is likely to be associated with less uncertainty in consumption and a reduced need for precautionary saving. These two effects go in the same direction and are therefore potentially confounded in the analysis above.

However, a consideration of household savings may allow a distinction to be made. If the jump in expenditure is solely due to liquidity constraints, then one would expect the household to dig into any savings in the run-up to becoming eligible. In the absence of uncertainty, the household will as far as possible try to smooth consumption across the income discontinuity, and this is inconsistent with maintaining positive savings. Put differently, a liquidity constraint cannot truly bind as long as savings are positive. Pension take-up should therefore be associated with an increase, or at least not a decrease, in household savings.

On the other hand, assume that the jump in consumption is entirely due to precautionary saving. Entering a phase of greater consumption certainty reduces the need to keep a buffer of savings. Therefore, the prediction is a fall in savings. This suggests a seemingly straightforward method of distinguishing between the two mechanisms: If behaviour is primarily guided by liquidity constraint, savings would be flat (at 0) or increase as income increases. By contrast, if the expenditure response to increased

income is due to precautionary saving, then household savings should decrease when income increases.

However, the data impose some limitations on this exercise. Only binary information on savings is available, in the form of the household's response to the question "Does this household, or a household member, own any of the following financial assets?" Yes/No responses are provided to each of the following options: Bank savings accounts, stokvels (ROSCAs), pension plans or retirement annuities, unit trusts, stocks or shares, cash loans expected to be repaid, life insurance, or any other savings. It is therefore possible to construct an overall savings indicator as follows: The household is coded as having savings (of *some* form) if it answered 'Yes' to at least one of these options, and otherwise not.

Clearly the analysis is limited to identifying households who move from being savers to non-savers or vice versa. Households who adjust their amount of savings from one positive amount to another cannot be discerned from those who do not adjust their (positive) savings at all.

Another potential problem with the data is related to how the respondents understand the question, in particular the last option, 'Any other savings'. There may be a tendency not to count modest amounts of cash kept in a purse or wallet as savings, whereas the same amount held in a bank account may be thought of as such. The resulting measurement error may lead to under-reporting of savings for those households who do not use a bank account or other non-cash savings technologies. If households are less likely to hold their savings in cash after the increase in income, then an estimate of the effect of the income on savings may be biased upwards.

Notwithstanding the limitations, it is interesting to look at what happens to the constructed savings indicator as the households become eligible for the pension. The results of these regressions are reported in Table 3.5. Here, the dependent variable is the binary savings indicator, but otherwise the regressions closely match those reported in Tables 3 and 4. The precautionary saving story unambiguously predicts a reduction in savings, while the liquidity constraints hypothesis is related to either no change or an increase in household savings. It is therefore natural to take the liquidity constraints story as the null hypothesis in these regressions, to be rejected in favour of precautionary saving if the coefficient on pension take-up or eligibility is negative and significant.

In columns 1–3, pension take-up is associated with a significant increase in the propensity to have some savings, in the order of 6–7 %-points, even when household demographics and the presence of older household members are controlled for. Columns 4–6 are again reduced-form regressions, relating savings to pension age-eligibility. Column 4 finds a positive and significant effect of eligibility on savings. When household demographics are controlled for in columns 5 and 6, the coefficients of interest are small and not significantly different from zero. In summary, the savings variable seems to either stay constant or increase in response to increased income.

Overall, there is no evidence of a decline in savings, and the null hypothesis of liquidity constraints cannot be rejected in favour of precautionary saving.

3.4.5 Expenditure response to an anticipated reduction in income

Above it was found that household expenditure (even when de-trended) responds positively to an anticipated increase in income. Moreover, the response of household savings to the anticipated increase in income seems to point in the direction of liquidity constraints, rather than precautionary saving, as an explanation for the consumption behaviour. However, this conclusion still rests on the assumptions of the standard model from which the Euler equation was derived. In this section the aim is to contrast the general framework of forward-looking but liquidity-constrained households with the stark Keynesian alternative in which households simply consume a fixed fraction of their income.

More advanced versions are considered in the literature, but they share the critical feature that current consumption is a positive function of current income. Deaton (1992) considers several rules of thumb with this feature, of which his benchmark rule is 'consume all cash on hand up to mean income, and 30 per cent of any excess'.

A Keynesian model of consumption compares well to the standard model as far as the results so far are concerned: For Keynesian households, an increase in income is accompanied by an increase in consumption, even if the change is anticipated. Furthermore, if consumption is a positive proportion of income, and savings are the difference between income and consumption, then savings must also be a fixed proportion of income. Therefore, the basic Keynesian model predicts household savings to increase when income does, which is supported by the findings in the previous section.

But an anticipated *reduction* in income may help distinguish between the models: If consumers are liquidity-constrained but forward-looking, then an anticipated decrease in income should not have any effect on expenditure because the required technology to smooth consumption is savings, not credit. On the other hand, Keynesian consumption would predict expenditure to decrease in line with income.

It is important to be clear that this strategy to differentiate between liquidity constraints and Keynesian consumption relies on the presence of an effective savings technology. Anyone can save in cash at home or on their person, but the expected return on savings may be negative if there is considerable risk of loss, appropriation by household members or theft. Even during apartheid, the South African post office offered a savings account which was available to the whole population. But it is possible that distance or transport cost to the nearest branch, mistrust in the system or illiteracy may have effectively blocked off this opportunity to some potential savers. Since the end of apartheid, most if not all banks offer savings accounts to anyone, though high fees and fear of intimidation may still act as barriers. Informal savings devices, for example in the form stokvels (ROSCAs), remain popular amongst black

South Africans. Taken together, it seems reasonable to assume that savings technology of some form is available to most South African households, and certainly more readily accessible than credit.

The South African child support grant is available to 'primary carers' (predominantly, mothers) of young children. There is a means test which will be ignored in this study because most black households easily pass it. Instead the focus is once again on the age cut-offs. Originally the grant covered any number of biological children up to and including the age of 6. So as a first pass, one could look for a drop in household expenditure when children turn 7. If there is such a drop, it could be taken as evidence for Keynesian behaviour rather than credit constraints.

But this is not entirely satisfactory. In the previous section, identification relied partly on the assumption that households where a, say 55–59-year-old woman is present is not substantially different from a household with a 60–64-year-old woman, except that the latter household would be age-qualified for the old age pension. In the case of children, it is less clear that a child with a 6-year-old is comparable to a 7-year-old. Young children develop fast and progress through the education system, so arguably an age difference of a single year may be associated with considerable changes in household behaviour even if other circumstances were fixed. For this reason, a change in the age cut-off for the child support grant will be exploited here. From 1 April 2003, the upper age limit went up from 6 to 8 years. So in the first two rounds of the survey data, collected in September 2001 and September 2002, household with children up to and including 6 years were age-qualified for a child-support grant. But in the final round used here, conducted in September 2003, households with children up to and including 8 years were age-eligible.

If the household consumption pattern over time is determined by Keynesian behaviour as opposed to liquidity constraints, then expenditure for a household with a child aged 7 in 2003 should be higher than the expenditure for a household with a child aged 7 in 2002, because only the former household is eligible for the grant. To see whether this is the case, run a regression of the form

$$Y_{it} - Y_{i,t-1} = \alpha_i + \mu_t + \gamma C_{it} + \delta C_{it} * I_{t=2003} + X_{it}\beta + \varepsilon_{it}.$$

Note that consumption is again first-differenced to account for within-household consumption trends. Here, C_{it} denotes the presence of a child age 7 in household i and year t. The indicator $I_{t \in S}$ is 1 if $t \in S$ and 0 otherwise. If households are subject to constraints on liquidity but not on savings, the coefficient δ should be zero: There is no difference between the consumption growth behaviour of a household with a child who turns seven in 2002, and that of a household with a child who turns seven in 2003, apart from what can be explained by the macro time-trend or changes in other characteristics. But Keynesian households will reduce consumption in line with the lapse of the child-support grant. Therefore, a household with a 7-year-old in 2002

should decrease its consumption compared to the previous year, whereas a household with a 7-year-old in 2003 should not. In other terms, the coefficient δ is zero under the standard model with liquidity constraints, and positive under the Keynesian model.

Table 3.6 presents the results of this regression. In column 1 the regression is run without control variables apart from household and survey fixed effects. The coefficient of interest is negative but not significantly different from zero. In column 2 household demographics are added as control variables but the coefficient of interest changes little and remains insignificant. In column 3 basic further controls are added in the form of variables counting the number of children in age categories from 0 up to 9 years of age. The coefficient of interest remains virtually unchanged and still not significant.

Thus the coefficient is far from being significantly different from zero. The magnitude is also negative at around 0.1 for all three regressions, which goes against the Keynesian model. The upshot is therefore that the null hypothesis of a standard model with forward-looking households and liquidity constraints cannot be rejected in favour of the Keynesian model.

Though it is likely that the relevant population was reasonably well informed of the basic qualifying requirements for the child support grant, the above analysis implicitly assumes that the *changes* in the qualifying age were also known. The analysis does not permit a clear distinction between the behaviour of forward-looking but 'surprised' households who learnt that their 7-year-olds were still receiving the grant in 2003, from that of non-forward-looking Keynesian households. Both types of households would be expected to consume more than their counterparts in 2002 before the reform. The results presented here, a negative coefficient not significantly different from zero, are therefore weakly indicative of forward-looking households who were aware of the changes in the rules.

3.5 Discussion and conclusion

There are a number of caveats to the findings of this paper. One is that there is no direct data on consumption, so that household expenditure is used as a proxy here. But the two are not always the same. For instance, the purchase of a durable good often entails an up-front expenditure, but the good is typically 'consumed' over several periods. Likewise, debt repayment may be regarded as an expense. But the idea of credit is precisely to separate in time the outlay associated with a good, and its consumption.

It is also clearly a limitation that the expenditure data is discrete. Binning the data entails a loss of information. Arguably, this makes it harder for the econometric tests to detect changes in expenditure since only movements across the bin boundaries are observed. This will tend to make it harder to reject a null hypothesis than it would be if the data were scalar. In terms of the findings, this may mean that the rejection of

precautionary saving and Keynesian consumption are less robust than they could be if expenditure was measured with more precision. In the case of Keynesian consumption, this concern is accentuated by the fact the child support grant is quantitatively less important than the old age pension.

This paper has assumed that savings technology is readily available, but this may not be the case. Besley (1995*a*) provides a number of reasons why savings may not be an available or attractive way to achieve smooth consumption in a typical developing country setting. On the other hand, South Africa has a relatively well-developed formal financial system. Though it has been difficult for poor people (without formal employment) to obtain fully-fledged bank accounts, savings accounts are more easily available. In particular, the post office has a reasonable presence even in rural areas and their savings books were available to anyone even during the Apartheid era, and in the sample studied here a full 48% report having savings in some form.

Many economists are quick to point to liquidity constraints when they suspect failures of the traditional consumption model. Indeed, the assumption of constrained credit is a staple of economic development theory. Yet direct and rigorous tests of this assumption, of the type that has recently gained popularity in analyses of rich-country data, is scarce. Moreover, in many cases findings consistent with liquidity constraints are equally compatible with precautionary saving or Keynesian consumption, and few studies have attempted to distinguish between these in a developing country context. This paper tests and rejects the standard model with perfect capital markets using data on a panel of black South African households. It also provides evidence that liquidity constraints, rather than precautionary saving or Keynesian behaviour, drive the observed excess sensitivity of consumption to anticipated income changes.

But the significance of the findings presented here goes beyond a mere cementation of what is already widely assumed in the theory of development economics. Credit constraints have implications for welfare. At the micro-level, efficient credit markets allow consumption smoothing and boost income through efficient allocation of resources for investment. At the macro-level, these investment misallocations are widely believed to impede growth. Credit markets matter for development, and effective development policy must be sensitive to whether or not households face liquidity constraints.

Figure 3.1: Household take-up of old-age pension by age of oldest member

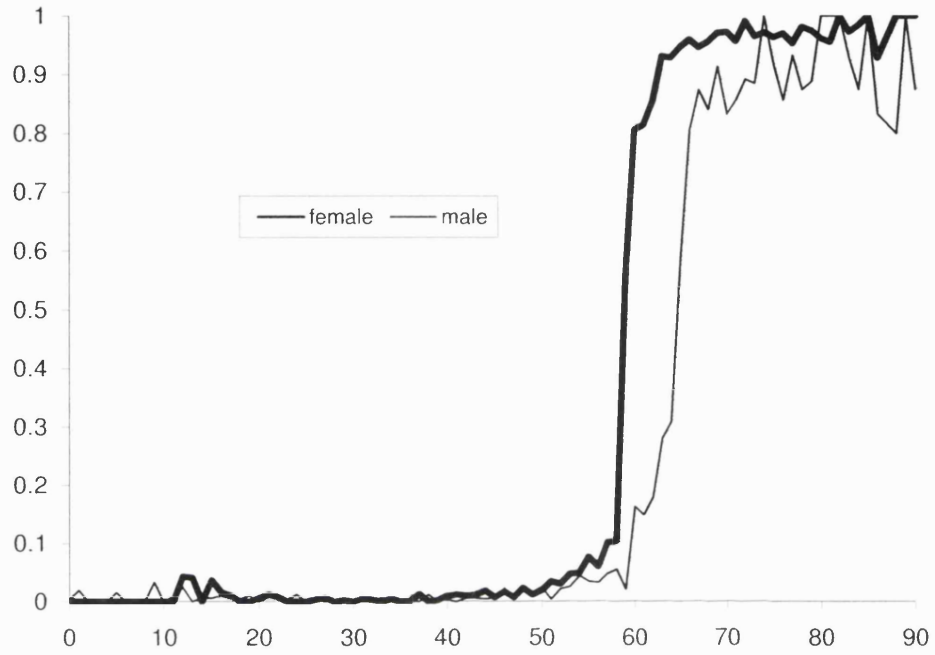

The thick line shows the fraction of households that receive the old age pension, as a function of the age of the oldest female household member. In order to focus on take-up as a function of female eligibility, households with men aged 55+ are excluded from this graph. Similarly, the thin line shows the household take-up rate of the pension as a function of the age of the oldest male household member. Households with women aged 55+ are excluded from the graph. Recall that women are age-eligible from the age of 60, and men from the age of 65. Age-eligibility is clearly a strong predictor of pension take-up. The data is from the September 2003 survey.

Table 3.1. Descriptive statistics.

	Mean	St. Dev.
Household size	3.91	2.74
Household receives the age-old pension	0.21	0.40
Household is age-eligible for pension	0.25	0.52
Household receives the child support grant	0.11	0.31
Log monthly expenditure	6.30	0.92
Household has savings in any form	0.48	0.50
Number of household-year observations	61,458	
Number of households observed once	17,577	
Number of households observed twice	11,964	
Number of households observed three times	6,651	
Total number of households	36,192	

Table 3.2. Household age-eligibility predicts pension take-up.

Dependent variable: Pension take-up (binary)

	1	2	3
Household is age-eligible	.487***	.392***	
	[.00904]	*[.0108]*	
Men 55+		-.0202*	0.00067
		[.0111]	*[.0117]*
Men 60+			.108***
			[.016]
Men 65+ (age-eligible)			.224***
			[.0219]
Men 70+			0.00557
			[.0188]
Women 55+		.116***	.0576***
		[.0101]	*[.0105]*
Women 60+ (age-eligible)			.413***
			[.0156]
Women 65+			.0981***
			[.0141]
Women 70+			.0693***
			[.0122]
Household demographics 0-54 years	No	Yes	Yes
Observations	61409	61409	61409
Households	36168	36168	36168

Linear regressions. The dependent variable is an indicator for whether the household receives the old age pension (1) or not (0). All regressions include household fixed effects. The household demographics not reported are variables for the number of female and male household members in each of the age groups 0-4, 5-14, 15-24, 25-34, 35-44 and 45-54, and household size. Robust standard errors, shown in brackets, are clustered at the level of the survey primary sampling unit. *** significant at 1%, ** significant at 5%, * significant at 10%.

Table 3.3. Expenditure response to pension income

Dependent variable: Household expenditure

	1	2	3	4	5	6
Household receives old-age pension	.195***	.126***	.107***			
	[.0199]	*[.0219]*	*[.0234]*			
Household is age-eligible for pension				.153***	.0825***	
				[.0157]	*[.0213]*	
Men 55+		.0633**	.0879***		.0518*	.087***
		[.0265]	*[.033]*		*[.0273]*	*[.033]*
Men 60+			-0.0578			-0.0457
			[.0373]			*[.0373]*
Men 65+ (age-eligible)			0.013			0.0374
			[.0405]			*[.0404]*
Men 70+			0.0553			0.0552
			[.0407]			*[.0407]*
Women 55+		-0.00431	-0.0309		-0.00847	-0.0247
		[.021]	*[.0252]*		*[.0234]*	*[.0252]*
Women 60+ (age-eligible)			.0584*			.103***
			[.0313]			*[.0305]*
Women 65+			0.0127			0.0229
			[.0295]			*[.0294]*
Women 70+			-0.0121			-0.00444
			[.0307]			*[.0307]*
Household demographics 0-54 years	No	Yes	Yes	No	Yes	Yes
Observations	60506	60506	60506	60532	60532	60532
Households	35877	35877	35877	35889	35889	35889

Linear regressions. The dependent variable is the logarithm of the imputed monthly household expenditure. All regressions include household fixed effects. The household demographics not reported are variables for the number of female and male household members in each of the age groups 0-4, 5-14, 15-24, 25-34, 35-44 and 45-54, and household size. Robust standard errors, shown in brackets, are clustered at the level of the survey primary sampling unit. *** significant at 1%, ** significant at 5%, * significant at 10%.

Table 3.4. Response in first-differenced expenditure to pension income

Dependent variable: Differenced household expenditure

	1	2	3	4	5	6
Household receives old-age pension	.326***	.231***	.226***			
	[.0632]	*[.0736]*	*[.0796]*			
Household is age-eligible for pension				.245***	.124*	
				[.05]	*[.0684]*	
Men 55+		.154*	0.106		0.14	0.0988
		[.0905]	*[.106]*		*[.0911]*	*[.106]*
Men 60+			0.0249			0.0523
			[.124]			*[.124]*
Men 65+ (age-eligible)			0.0567			0.107
			[.145]			*[.143]*
Men 70+			0.0477			0.0548
			[.148]			*[.149]*
Women 55+		.149**	.158**		.162**	.176**
		[.0662]	*[.0764]*		*[.0714]*	*[.076]*
Women 60+ (age-eligible)			0.0203			0.109
			[.102]			*[.1]*
Women 65+			-0.0132			0.0105
			[.0972]			*[.0972]*
Women 70+			-0.0732			-0.0561
			[.0888]			*[.0888]*
Household demographics 0-54 years	No	Yes	Yes	No	Yes	Yes
Observations	24119	24119	24119	24123	24123	24123
Households	17739	17739	17739	17742	17742	17742

Linear regressions. The dependent variable is the first-differenced logarithm of the imputed monthly household expenditure. All regressions include household fixed effects. The household demographics not reported are variables for the number of female and male household members in each of the age groups 0-4, 5-14, 15-24, 25-34, 35-44 and 45-54, and household size. Robust standard errors, shown in brackets, are clustered at the level of the survey primary sampling unit. *** significant at 1%, ** significant at 5%, * significant at 10%.

Table 3.5. Response in household savings to pension income

Dependent variable: Household has savings in any form (binary indicator)

	1	2	3	4	5	6
Household receives old-age pension	.07***	.0592***	.0644***			
	[.0132]	*[.0144]*	*[.0159]*			
Household is age-eligible for pension				.0309***	0.0101	
				[.0107]	*[.0133]*	
Men 55+		0.0156	0.0308		0.016	0.0298
		[.017]	*[.0205]*		*[.0173]*	*[.0204]*
Men 60+			-0.0222			-0.0154
			[.0222]			*[.0221]*
Men 65+ (age-eligible)			-0.0408			-0.0278
			[.0269]			*[.0267]*
Men 70+			.0507*			.0512*
			[.0272]			*[.0273]*
Women 55+		-0.021	-0.0189		-0.00692	-0.0157
		[.0136]	*[.0157]*		*[.0145]*	*[.0157]*
Women 60+ (age-eligible)			-0.00998			0.0169
			[.0209]			*[.0203]*
Women 65+			0.012			0.0183
			[.0207]			*[.0206]*
Women 70+			-0.00621			-0.00174
			[.0217]			*[.0216]*
Household demographics 0-54 years	No	Yes	Yes	No	Yes	Yes
Observations	61409	61409	61409	61458	61458	61458
Households	36168	36168	36168	36192	36192	36192

Linear regressions. The dependent variable is 1 if the household has any savings of any form, and 0 otherwise. All regressions include household fixed effects. The household demographics not reported are variables for the number of female and male household members in each of the age groups 0-4, 5-14, 15-24, 25-34, 35-44 and 45-54, and household size. Robust standard errors, shown in brackets, are clustered at the level of the survey primary sampling unit. *** significant at 1%, ** significant at 5%, * significant at 10%.

Table 3.6. Response in first-differenced expenditure to the foreseeable lapse of the child support grant

Dependent variable: Differenced household expenditure	1	2	3
Presence of child aged 7	0.092 *[.0593]*	0.0546 *[.0602]*	0.0552 *[.0717]*
Presence of child aged 7 in 2003	-0.111 *[.0824]*	-0.0975 *[.0821]*	-0.0997 *[.0869]*
Household and survey fixed effects	*Yes*	*Yes*	*Yes*
Household demographics	*No*	*Yes*	*Yes*
Controls for number of children aged 0, 1, 2, ... , 9	*No*	*No*	*Yes*
Observations	24123	24123	24123
Households	17742	17742	17742

Linear regression of the first-differenced logarithm of imputed household expenditure on a binary variable indicating the presence in the household of a child aged 7, and and indicator for the presence of a child aged 7 in 2003. The coefficient on the latter should be zero under the null hypothesis of forward-looking households and liquidity constraints. The null hypothesis is not rejected. The household demographics are variables for the number of female and male household members in each of the age groups 0-4, 5-14, 15-24, 25-34, 35-44 and 45-54, 55+ as well as household size. The controls for the number of children are count variables for children of age 0, 1, ..., 9 in the household. Robust standard errors, shown in brackets, are clustered at the level of the survey primary sampling unit. *** significant at 1%, ** significant at 5%, * significant at 10%.

References

Adamolekun, K. 2001. "Survivors' motives for extravagant funerals among the Yorubas of Western Nigeria." *Death Studies* 25(7):609–619.

Altonji, J. G. & A. Siow. 1987. "Testing the response of consumption to income changes with (noisy) panel data." *Quarterly Journal of Economics* 102(2):293–328.

Anderson, S. & J. M. Baland. 2002. "The Economics of ROSCAs and Intrahousehold Allocation." *Quarterly Journal of Economics* 117(3):963–95.

Ardington, C., D. Lam, M. Leibbrandt & J. Levinsohn. 2004. "Savings, Insurance, and Debt over the Post-Apartheid Experience: A Review of Recent Research." *South African Journal of Economics* 72(3):621–657.

Ardington, C. & M. V. Leibbrandt. 2004. "Financial Services and the Informal Economy." Working Paper, Centre for Social Science Research, University of Cape Town.

Arhin, K. 1994. "The Economic Implications of Transformations in Akan Funeral Rites." *Africa: Journal of the International African Institute* 64(3):307–322.

Arnott, R. & J. E. Stiglitz. 1991. "Moral Hazard and Nonmarket Institutions - Dysfunctional Crowding out or Peer Monitoring." *American Economic Review* 81(1):179–190.

Aryeetey, E. & C. Udry. 2000. "Saving in Sub-Saharan Africa." CID Working Papers, Center for International Development at Harvard University.

Ashraf, N., D. Karlan & W. Yin. 2006*a*. "Tying Odysseus to the mast: Evidence from a commitment savings product in the Philippines." *Quarterly Journal of Economics* 121(2):635–672.

Ashraf, N., D. S. Karlan & W. Yin. 2006*b*. Household Decision Making and Savings Impacts: Further Evidence from a Commitment Savings Product in the Philippines. Center Discussion Paper 939 Yale Economic Growth Center.

Attanasio, O. P., L. Picci & A. E. Scorcu. 2000. "Saving, Growth, and Investment: A Macroeconomic Analysis Using a Panel of Countries." *The Review of Economics and Statistics* 82(2):182–211.

Banerjee, A. & E. Duflo. 2005. Growth Theory through the Lens of Economic Development. In *Handbook of Economic Growth*, ed. P. Aghion & S. N. Durlauf. Vol. 1A Amsterdam: North-Holland chapter 7.

Banerjee, A. V. 2003. "Contracting Constraints, Credit Markets, and Economic Development." *Advances in Economics and Econometrics: Theory and Applications: Eighth World Congress of the Econometric Society* .

Banerjee, A. V. & E. Duflo. 2008. "Do Firms Want to Borrow More? Testing Credit Constraints Using a Directed Lending Program." Working paper.

Barham, B. L., S. Boucher & M. R. Carter. 1996. "Credit constraints, credit unions, and small-scale producers in Guatemala." *World Development* 24(5):793–806.

Bernheim, B. D. 1991. "How Strong Are Bequest Motives? Evidence Based on Estimates of the Demand for Life Insurance and Annuities." *Journal of Political Economy* 99(5):899–927.

Bertrand, M., S. Mullainathan & D. Miller. 2003. "Public Policy and Extended Families: Evidence from Pensions in South Africa." *The World Bank Economic Review* 17(1):27–50.

Bertrand, M., S. Mullainathan & E. Shafir. 2004. "A Behavioral-Economics View of Poverty." *American Economic Review* 94(2):419–423.

Besley, T. 1995*a*. "Nonmarket Institutions for Credit and Risk Sharing in Low-Income Countries." *The Journal of Economic Perspectives* 9(3):115–127.

Besley, T. 1995*b*. Savings, Credit and Insurance. In *Handbook of Development Economics*, ed. J. Berhman & T.N. Srinivasan. Vol. 3a Amsterdam: North-Holland chapter 36.

Besley, T. & C. Meghir. 1998. "Tax Based Savings Incentives." Unpublished manuscript.

Besley, T., S. Coate & G. Loury. 1993. "The Economics of Rotating Savings and Credit Associations." *The American Economic Review* 83(4):792–810.

Bloch, F., G. Genicot & D. Ray. 2008. "Informal insurance in social networks." *Journal of Economic Theory* 143(1):36–58.

Bold, T. 2007. "Risk-Sharing in Insurance Groups in Rural Ethiopia." Work in Progress.

Browning, M. 2000. "The Saving Behaviour of a Two-Person Household." *The Scandinavian Journal of Economics* 102(2):235–251.

Browning, M. & A. Lusardi. 1996. "Household Saving: Micro Theories and Micro Facts." *Journal of Economic Literature* 34(4):1797–1855.

Browning, M. & M. D. Collado. 2001. "The Response of Expenditures to Anticipated Income Changes: Panel Data Estimates." *American Economic Review* 91(3):681–692.

Bryant, J. & A. Prohmmo. 2002. "Equal contributions and unequal risks in a North-east Thai village funeral society." *The Journal of Development Studies* 38(3):63–75.

Budlender, D., S. Rosa & K. Hall. 2005. "At all costs? Applying the means test for the Child Support Grant." Children's Institute and Centre for Actuarial Research, University of Cape Town.

Case, A. & A. Deaton. 1998. "Large cash transfers to the elderly in South Africa." *Economic Journal* 108(450):1330–1361.

Case, A., A. Garrib, A. Menendez & A. Olgiati. 2008. "Paying the Piper: The High Cost of Funerals in South Africa." NBER Working Paper 14456.

Case, A., V. Hosegood & F. Lund. 2005. "The Reach and Impact of Child Support Grants: Evidence from KwaZulu-Natal." *Development Southern Africa* 22(4):467–82.

Claessens, S. 2006. "Access to Financial Services: A Review of the Issues and Public Policy Objectives." *The World Bank Research Observer* 21(2):207–240.

Collier, P. & J. W. Gunning. 1999. "Explaining African Economic Performance." *Journal of Economic Literature* 37(1):64–111.

Cordery, S. 2003. *British Friendly Societies, 1750-1914.* Basingstoke: Palgrave MacMillan.

de Witte, M. 2003. "Money and death: Funeral business in Asante, Ghana." *Africa: Journal of the International African Institute* 73(4):531–559.

Deaton, A. 1991. "Saving and Liquidity Constraints." *Econometrica* 59(5):1221–48.

Deaton, A. 1992. "Household saving in LDCs: Credit markets, insurance and welfare." *The Scandinavian Journal of Economics* 94(2):253–273.

Dercon, S. 2002. "Income Risk, Coping Strategies, and Safety Nets." *The World Bank Research Observer* 17(2):141–166.

Dercon, S., J. De Weerdt, T. Bold, & A. Pankhurst. 2006. "Group-based Funeral Insurance in Ethiopia and Tanzania." *World Development* 34(4):685–703.

Dhemba, J., P. Gumbo & J. Nyamusara. 2002. "Social Security in Zimbabwe: Phase II: Zunde raMambo and Burial Societies." *Journal of Social Development in Africa* 17(2):132–153.

Duflo, E. 2003. "Grandmothers and Granddaughters: Old-Age Pensions and Intrahousehold Allocation in South Africa." *The World Bank Economic Review* 17(1):1–25.

Duflo, E. 2006. Poor but Rational? In *Understanding Poverty*, ed. A. Banerjee, R. Bénabou & D. Mookherjee. New York: Oxford University Press pp. 367–78.

Duflo, E. & C. R. Udry. 2003. "Intrahousehold Resource Allocation in Côte d'Ivoire: Social Norms, Separate Accounts and Consumption Choices." NBER Working Paper.

Edmonds, E. V. 2006. "Child labor and schooling responses to anticipated income in South Africa." *Journal of Development Economics* 81(2):386–414.

Engen, E. M., W. G. Gale, J. K. Scholz, B. D. Bernheim & J. Slemrod. 1994. "Do Saving Incentives Work?" *Brookings Papers on Economic Activity* 1994(1):85–180.

Fafchamps, M. & S. Lund. 2003. "Risk-sharing networks in rural Philippines." *Journal of Development Economics* 71:261–287.

Fingland Jack, A. 1912. *An Introduction to the History of Life Assurance.* London: P. S. King & Son.

FinMark Trust. 2003. "Summary Report: Findings of the FinScope Study Into Financial Access and Behavior of the South African Population 2003." Brochure.

Fischer, S. 1973. "A Life Cycle Model of Life Insurance Purchases." *International Economic Review* 14(1):132–152.

Flavin, M. 1985. "Excess sensitivity of consumption to current income: liquidity constraints or myopia?" *Canadian Journal of Economics* 18(1):117–136.

Gersovitz, M. 1988. Saving and development. In *Handbook of Development Economics*, ed. Hollis Chenery & T.N. Srinivasan. Vol. 1 Amsterdam: North-Holland chapter 10.

Gugerty, M. K. 2007. "You Can't Save Alone: Commitment and Rotating Savings and Credit Associations in Kenya." *Economic Development and Cultural Change* 55:251–282.

Hall, R. E. 1978. "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence." *The Journal of Political Economy* 86(6):971–987.

Harris, C. & D. Laibson. 2003. Hyberbolic Discounting and Consumption. In *Advances in Economics and Econometrics: Theory and Applications: Eighth World Congress.* Cambridge: Cambridge University Press.

Heckman, J. 1997. "Instrumental Variables: A Study of Implicit Behavioral Assumptions Used in Making Program Evaluations." *The Journal of Human Resources* 32(3):441–462.

Hertz, T. 2005. "The Effect of Minimum Wages on the Employment and Earnings of South Africa's Domestic Service Workers." Upjohn Institute Working Paper 05-120.

Hsieh, C. T. 2003. "Do Consumers React to Anticipated Income Changes? Evidence from the Alaska Permanent Fund." *American Economic Review* 93(1):397–405.

Johnson, P. 1985. *Saving and spending: The Working-Class Economy in Britain, 1870-1939.* Oxford: Oxford University Press.

Johnston, H. W. 1903. *The Private Life of the Romans.* Chicago: Scott, Foresman and Co. chapter 14.

Keswell, M. 2004. "Social Networks, Extended Families, and Consumption Smoothing: Field Evidence from South Africa." Technical Report, University of Cape Town.

King, R. G. & R. Levine. 1993. "Finance and Growth: Schumpeter Might be Right." *The Quarterly Journal of Economics* 108(3):717–737.

Kirby, K. N. & N. N. Marakovic. 1996. "Delay-Discounting Probabilistic Rewards: Rates Decrease as Amounts Increase." *Psychonomic Bulletin & Review* 3(1):100–104.

Levenson, A. R. & T. Besley. 1996. "The anatomy of an informal financial market: Rosca participation in Taiwan." *Journal of Development Economics* 51(1):45–68.

Lewis, F. D. 1989. "Dependents and the Demand for Life Insurance." *The American Economic Review* 79(3):452–467.

Ligon, E., J. P. Thomas & T. Worrall. 2002. "Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies." *Review of Economic Studies* 69(1):209–244.

Mankiw, N. G. 1982. "Hall's Consumption Hypothesis and Durable Goods." *Journal of Monetary Economics* 10(3):417–425.

Morduch, J. 1990. "Risk, Production and Saving: Theory and Evidence from Indian Households." Manuscript, Harvard University.

Morduch, J. 1999. "Between the State and the Market: Can Informal Insurance Patch the Safety Net?" *The World Bank Research Observer* 14(2):187–207.

National Treasury. 2005. "Provincial Budgets and Expenditure Review.".

Ngwenya, B.N. 2003. "Redefining kin and family social relations: burial societies and emergency relief in Botswana." *Journal of Social Development in Africa* 18(1):85–110.

Parker, J. A. 1997. "The Reaction of Household Consumption to Predictable Changes in Payroll Tax Rates." Working Papers 9724 Wisconsin Madison - Social Systems.

Parrott, A. L. 1985. *The Iron Road to Social Security.* Lewes, Sussex: Book Guild Ltd.

Paxson, C. 1992. "Using Weather Variability to Estimate the Response of Savings to Transitory Income in Thailand." *American Economic Review* 82(1):15–33.

Porteous, D. J. & E. Hazelhurst. 2004. *Banking on Change.* Cape Town: Double Storey.

Poterba, J. M., S. F. Venti & D. A. Wise. 1996. "Personal Retirement Saving Programs and Asset Accumulation: Reconciling the Evidence." NBER working paper 5599.

Quisumbing, A. R. & J. A. Maluccio. 2003. "Resources at Marriage and Intrahousehold Allocation: Evidence from Bangladesh, Ethiopia, Indonesia, and South Africa." *Oxford Bulletin of Economics and Statistics* 65(3):283–327.

Ranchhod, V. 2006. "The Effect of the South African Old Age Pension on Labour Supply of the Elderly." *South African Journal of Economics* 74(4):725–744.

Rasmussen, K. 2002. "Securing Consumption in Low-Income Households: An Empirical Study of Coping with Shocks Amongst Black Households in the KwaZulu-Natal Province of South Africa 1993–1998." Unpublished Masters thesis, University of Copenhagen, Denmark.

Ray, D. 2008. Development Economics. In *The New Palgrave Dictionary of Economics*, ed. S. N. Durlauf & L. E. Blume. Second ed. Palgrave Macmillan.

Rosa, S. & C. Mpokotho. 2004. "Extension of the child support grant to children under 14 years: Monitoring Report." Children's Institute, University of Cape Town.

Rosenzweig, M. R. & K. I. Wolpin. 1993. "Credit Market Constraints, Consumption Smoothing, and the Accumulation of Durable Production Assets in Low-Income Countries: Investments in Bullocks in India." *The Journal of Political Economy* 101(2):223–244.

Roth, J. 2001. "Informal microinsurance schemes - the case of funeral insurance in South Africa." *Small Enterprise Development* 12(1):39–50.

Rutherford, Stuart. 2000. *The Poor and Their Money.* New Delhi: Oxford University Press.

Statistics South Africa. 2005. "Labour Force Survey September 2005." Statistical release P0210.

Stegman, M. A. 1999. *Savings for the Poor: The Hidden Benefits of Electronic Banking.* Washington, DC: Brookings Institution Press.

Strotz, R. H. 1956. "Myopia and Consistency in Dynamic Utility Maximization." *The Review of Economic Studies* 23(3):165–180.

Thomson, R. J. & D. B. Posel. 2002. "The Management of Risk by Burial Societies in South Africa." *South African Actuarial Journal* 2:83–127.

Townsend, R. M. 1994. "Risk and Insurance in Village India." *Econometrica* 62(3):539–591.

Trenerry, C. F. 1926. *The Origin and Early History of Insurance: Including the Contract of Bottomry.* London: P. S. King & Son.

Udry, C. 1994. "Risk and Insurance in a Rural Credit Market: An Empirical Investigation in Northern Nigeria." *Review of Economic Studies* 61(3):495–526.

Van der Linden, M., ed. 1996. *Social Security Mutualism: The Comparative History of Mutual Benefit Societies.* Bern: Peter Lang.

Verhoef, G. 2001. "Informal Financial Service Institutions for Survival: African Women and Stokvels in Urban South Africa, 1930–1998." *Enterprise and Society* 2(2):259–296.

Vermeulen, F. 2002. "Collective Household Models: Principles and Main Results." *Journal of Economic Surveys* 16(4):533–564.

Warnecke, W. G. G. 1994. "Burial Societies, African Religion and the Church." Unpublished Mth dissertation, University of Natal.

Yaari, M. E. 1965. "Uncertain Lifetime, Life Insurance, and the Theory of the Consumer." *Review of Economic Studies* 32(2):137–150.