

Chapter 6

Inequality traps and human-capital accumulation in South Africa⁷⁰

Miquel Pellicer and Vimal Ranchhod

Introduction

A large body of theoretical literature has emerged in recent decades which explores the links between the degree of inequality in a society and its investment levels and patterns. This literature has the potential to yield useful insights for South Africa. South Africa has one of the highest levels of inequality in the world, as well as very high unemployment levels. Moreover, growth has been relatively sluggish in recent decades, in comparison to other emerging markets. The study of inequality and investment in the aforementioned literature has the potential to shed light on the conditions under which situations of high inequality, low growth and high unemployment can emerge. The key contributions of this literature are the acknowledgement that all of these variables are interrelated, as well as the focus on their interactions. In this chapter we examine the insights of this literature that are potentially most relevant to South Africa and tentatively explore the empirical relevance of the mechanisms emphasised. We believe that this exercise can be helpful in informing policy discussions.

There are several types of investments through which inequality and employment interact to produce poor economic outcomes. An important one is entrepreneurship. Inequality may limit entrepreneurship in certain contexts and thus limit employment levels.⁷¹ This would be a rather direct connection between inequality and employment, through the demand for labour. Another potential channel is via investment in innovation. Inequality may encourage or discourage innovation, thus affecting growth and employment levels.⁷² Attempting to account for all of the ways in which inequality and investment interact in a

70 We would like to thank Murray Leibbrandt, Anthony Black and the participants of the SANPAD Employment-Intensive Growth workshop for very useful comments and suggestions. All errors are our own.

71 See, for instance, Banerjee and Newman (1993).

72 See Foellmi and Zweimüller (2006).

coherent and compact way would be unmanageable. Thus, in order to narrow down the problem, we ignore many of these channels and focus specifically on one: investment in education. We believe that education differentials are a crucial element for understanding South Africa's inequality, unemployment and poor growth.

Regarding inequality, recent empirical evidence shows that labour-market differentials account for most of the observed income inequality in South Africa (Leibbrandt et al., 2010). Wage differentials between skilled and unskilled labour in South Africa have likewise been shown to be extremely high compared to other countries (Lam, 1999).

With regard to unemployment, we present strong evidence later in this chapter that the unemployment rate of graduates with post-secondary qualifications is substantially lower than the average. From a partial equilibrium point of view, this implies that broadening access to higher education would foster employment. The effects can actually be substantially stronger after considering any general equilibrium effects: increasing the supply of skills is likely to make the remaining unskilled workers more desirable from the employers' point of view, thus relieving the unemployment problem of the unskilled. Finally, with regard to growth, skill accumulation would naturally contribute directly to growth via increased productivity, as well as indirectly through the effect mentioned above on unemployment. Thus, broadening access to higher education can be thought of as a labour-supply-side mechanism for achieving overall growth and, more particularly, employment growth.

In this context, we discuss the role of inequality in enabling or constraining employment and growth via the interactions of inequality and educational decisions. We discuss three different theoretical frameworks to study the interactions between inequality and educational choices. The following frameworks are ordered in terms of their complexity and (arguably) realism:

1. That markets work perfectly and that the social environment plays no role in decision-making.
2. That credit markets are imperfect and play no role in the social environment.
3. That the social environment affects information, preferences and actions.

These three frameworks lead to very different conclusions with regard to the effects of inequality on educational decisions, and to very different policy implications. According to the perfect-markets framework, inequality can actually be beneficial for growth and investment, as it generates incentives to obtain higher levels of education. In such a framework, there is no room for policy on a normative or on a positive basis. In the second framework of imperfect markets, inequality acts mainly as a constraint to investment. Inequality traps can emerge where high levels of inequality lead to little educational investment, which in turn regenerates

the high levels of inequality. In such a situation, the role of inequality for incentives is minimal, and is dwarfed by its constraining role. Implications for policy include individual-income redistribution and education financing, as well as, interestingly, the expansion of educational opportunities for intermediate levels of skill. Finally, the framework which takes social interactions into account leads to additional mechanisms through which inequality traps might arise: disadvantaged social environments may limit the aspirations or discourage the development of the abilities of their members, leading them to remain disadvantaged. This framework warns against the utility of individual-income redistribution and instead suggests policy interventions that target groups as a whole.

We then make use of two, large datasets to identify which of the hypotheses seems to have greater empirical support. The National Income Dynamics Study of 2008 is used to investigate returns to education for youth, educational attainment and the reasons for not continuing with one's education. All of this is done for different quintiles of household per capita income, which is how we operationalise inequality in our data. We then make use of the Cape Area Panel Study to identify some aspects relating to the third hypothesis, namely those dynamics driven by social influences. In particular, we look at young people's career aspirations, as well as how measures of scholastic aptitude vary with age for youth from different socio-economic backgrounds.

To briefly summarise our findings, we find most support for the second theoretical framework discussed. Youth in South Africa from all backgrounds would like to study further to obtain a post-secondary qualification, but income and credit constraints bind many of them. In addition, youth from poorer backgrounds are disadvantaged by their primary and secondary schooling experiences, to the extent that many of them do not graduate from school with sufficient grades to even be eligible for enrolment in a post-secondary qualification.

We conclude with a discussion of policy choices to improve education and reduce inequality in the South African context. Two types of objectives are considered: improving access to tertiary education and reducing differences in education quality at the primary and secondary levels. On the first, we argue that there is scope for improving access to information on existing education subsidy programmes, as well as increasing the coverage of these programmes to include registration fees. Moreover, there would be potential benefits and scope to improve access at lower levels of tertiary education. Regarding pre-tertiary education, we argue that a policy to integrate students from poorer backgrounds into high quality schools is likely to be fraught with problems and, at best, have a minor effect. In contrast, efforts focused on improving the quality of schools in disadvantaged communities would probably be more productive.

Inequality and investment: Insights from the theoretical literature

Introduction and general framework

A growing body of theoretical literature in economics has addressed the interactions between the level of inequality in a society and the investment choices of individuals within that society. Overall inequality and individual investment choices can feed into each other in potentially complex ways. There have been different approaches, or theories, that emphasise different aspects of these interactions, reaching significantly different conclusions as to the role of inequality for investment. The purpose of this section is to give a streamlined and non-technical account of these different approaches, focusing particularly on the decision to invest in education.⁷³

We shall discuss the different approaches to analysing interactions between inequality and individual choices using a common overall framework. The key components of this framework are the following. First, when facing a series of choices, individuals make the choice that is best for them, given their aspirations and values, their perceptions and their constraints. The different approaches below make different assumptions with regard to how aspirations, perceptions and constraints are formed, but they all assume that once they are taken into account, individuals try to make the best decision possible for themselves.

Second, we focus particularly on the choice of whether to obtain more/better education or not, although many of the insights would apply equally well to the decision to become an entrepreneur, for example. For concreteness, we shall often frame our analysis in terms of the decision to attend college or not. However, all relevant insights apply to other educational decisions, such as whether to attend better-quality education at primary and secondary levels.

We frame the decision of whether to obtain a tertiary education or not (or whether to attend higher-quality primary and secondary education or not) as entailing costs and benefits. The benefits are mainly the perceived improvement of labour-market prospects and economic status gained by improving one's education. These benefits may differ from person to person, first because different people have different abilities that make them exploit their educational investment differently and, second, because even people with similar abilities may hold different perceptions as to how the investment is rewarded in the society. In turn, the value attached to these perceived benefits may also differ from person to person, as people differ in aspirations and values. Some individuals may value a

73 For excellent technical surveys of many of the models discussed in this section, see Piketty (2000); Bertola (2000); Bertola et al. (2006).

higher economic status or fear the possibility of a low socio-economic status more strongly than others.

The educational investment also has costs. Here we focus particularly on monetary costs arising from tuition fees and living expenses, as well as potential psychic costs associated with the effort of studying. In addition to the direct benefits and costs of educational investments, we account for the fact that people may face constraints when making their decisions. These constraints may prevent people from undertaking their preferred option on the basis of cost–benefit calculations. We focus particularly on monetary constraints: individuals need to acquire the funds necessary to pay for the tuition and the living expenses while studying.

Thus, in our general framework of individual decisions, people have perceptions about the payoffs from education; they have aspirations that make them value these payoffs more or less; they consider their costs and constraints, and if the benefits outweigh the costs while satisfying their constraints, they then choose to obtain more education. This is a very stylised account of how actual decisions work. However, we can still capture what we believe are key elements in the decision-making process, and the framework is sufficiently flexible in the sense that it accommodates vastly different approaches to investment decisions.

Since we are interested in possible interactions between inequality and individual choices, we need to have a framework for how inequality is generated. For simplicity, and to focus as cleanly as possible on certain key mechanisms, we consider inequality as emerging primarily from the differences in pay and economic status that come from differences in productivity, which in turn derive from differences in education and ability.

This stylised framework allows us to discuss different, important approaches to the interactions between inequality and educational choices. We consider three approaches. We label the first the ‘perfect world’ because it abstracts from many real-world factors. Nonetheless, it is useful in the sense that it serves as a benchmark and because some policy debates are informed by its logic. In the ‘perfect-world’ approach, everyone knows the returns from education, which depend only on innate ability, and markets work perfectly. The second approach is the ‘capital-market-imperfections’ approach. This approach is the same as the previous one except for the fact that markets, and in particular the market for borrowing and lending, do not work perfectly. The third approach emphasises social interactions and social externalities. This is arguably the most realistic approach where aspirations, perceptions and the development of abilities depend on one’s environment, be it one’s neighbourhood or one’s social group, however defined.

Inequality and investment in a 'perfect world'

Consider a setting in which young individuals differ in their innate abilities and in their family background: some come from wealthy families and others from poor families. Suppose, moreover, that everyone knows the returns from education, which depend only on ability, and that tuition costs are the same for everyone. Who would like to go to college (or, more generally, invest in more/better education) in such a setting?

The benefits of education in such a setting would be higher for those with higher abilities and higher aspirations. The costs would be the same for everyone. Thus, only children with high abilities and high aspirations would choose to go to college, while for the rest, the costs would outweigh the benefits and they would prefer not to go to college.

How does this change when we consider wealth constraints? Would poor, bright children be able to go to college if they wished to? In a setting where markets work perfectly, the answer is yes. Poor, bright children would have good prospects in the labour market—they are very likely to end up with high earnings in the future. For them, the educational investment is very profitable and everyone knows this in our setting. Thus, for banks, financing such an investment could generate considerable returns. Banks, in turn, could obtain the necessary resources from rich parents with children of low ability. For these parents, the returns from the bank would be higher than the returns from the education of their children. The key point is that in a perfect world, what matters for investment is not the resources at one's disposal, but the potential benefits of the investment. What matters is not whether one is rich or poor, but whether one is bright or not. Sufficiently bright children can always find financing for their studies.

The fact that wealth does not affect investment choices, in turn, implies that inequality does not matter for access to college. Regardless of the level of inequality in a country, how many poor people there are or how poor they are, all bright people, whether they are rich or poor, would find it desirable to go to college and would be able to finance it. To be sure, in a very poor country, there would be relatively few funds available and borrowing would be costly. But then again, this implies that saving would be very rewarding and, ultimately, only the very brightest children would find it worthwhile going to college. Again, individual income would not matter for investment decisions and nor would the level of inequality.

In this setting, there is little scope for policy, neither normatively nor on efficiency grounds. From a normative point of view, wealth results from one's contribution to the economy in terms of the higher productivity afforded by ability and education. This distributional criterion could be considered fair. Moreover, to the extent that ability is not transmitted across generations, the resulting society

would display high levels of mobility. Even if ability were transmitted across generations, and poor and rich dynasties emerged, each dynasty would be rich or poor on the basis of ability, not on the basis of their prior wealth.

From an efficiency point of view, there are two types of policies that might be considered: (1) a subsidy to pay for tuition fees, and (2) progressive redistribution from rich to poor. An education subsidy would induce middle-ability individuals to study. However, it is not clear whether that would be a wise policy to follow in this case. In such an economy, there are no ‘unexploited educational opportunities’. Those that don’t go to college, choose not to because it would give them relatively low returns. Presumably, the resources that would go to the education subsidy could be put to better use, that is, invested in an activity that yields higher returns. Regarding the policy of redistributing income from rich to poor, redistribution towards poor dynasties would not affect education choices and would not have long-lasting effects. This would be the case even if ability were transmitted across dynasties, which would result in some inequality persistence. In fact, to the extent that such redistribution took resources from the educated and rich, it would reduce the benefits of education, and would actually drive the middle-income/middle-ability people away from colleges.

Inequality and investment with market imperfections

A large body of theoretical work has addressed the interactions between investment and individual investment choices in a context of imperfect capital markets.⁷⁴ Capital-market imperfections imply that certain individuals have limited access to credit. In such a situation, family wealth matters for investment choices: only individuals who are rich enough can afford to go to college. This leads to interesting and important interactions between the level of inequality and overall educational attainment, and leaves room for policy to improve economic outcomes.

The first question to address in settings of capital-market imperfections is: Why would capital markets not work properly in the first place? We focus here on the market for borrowing and lending. Borrowing and lending, and more generally any financial transaction, has the peculiarity that the different elements of the transaction occur at different moments in time. Transactions in traditional markets essentially occur simultaneously: when buying groceries, the product and the payment for the product get exchanged at the same time. When borrowing a

74 See, among others, Galor and Zeira (1993); Banerjee and Newman (1993); Aghion and Bolton (1997); Piketty (1997); Owen and Weil (1998); Maoz and Moav (1999); Ghatak and Jiang (2002) and Mokherjee and Ray (2003). The *World Development Report, 2006* by the World Bank offers an excellent non-technical survey of some of this literature.

sum of money, the product (that is the loan), is given first, and the payment (the reimbursement of the loan) is done later, sometimes after a substantial amount of time. This generates the possibility that the payment side of the transaction might not be honoured. In the case of educational investments, individuals may ask for an educational loan and then run away, or they may not exert much effort during the studies and end up unskilled, earning insufficient amounts to pay back the debt. The incentives to do such things increase with the amount of money that needs to be paid back. The higher the debt burden, the more attractive it becomes to rescind on one's obligations, and the less reward there is for one's efforts. For this reason, lenders cannot respond to default risks by increasing interest rates on loans, as this just worsens the problem. The solution for banks is to fund individuals that need to borrow only limited amounts, for whom the incentives to default are smaller. These individuals are the richer ones, so the end result is that the poor get excluded from credit.

In a setting of imperfect markets, therefore, wealth constraints matter. All high-ability individuals, be they poor or rich, would still wish to invest in education. But the poorer ones would not be able to do so because they would not obtain the necessary funding. In this context, the level of inequality in the society matters for overall access to education. A highly unequal society, in which there are many poor families and these are very poor, would suffer from low levels of education, as most of their population would be unable to finance it.

In this framework, inequality and education levels interact in meaningful ways. Consider, as seems sensible and realistic, that the relative wages of skilled and unskilled workers depend on the relative scarcity of skills: in an economy with severe skills scarcities, the few skilled individuals would command a strong premium while unskilled labour, being so abundant, would suffer from high levels of unemployment and low pay. This generates a potential feedback loop between inequality and education levels. High levels of inequality would generate few educated individuals, making skills scarce. This, in turn, would lead to high differences in pay between the educated (and rich) and the less educated (and poor); in other words, high inequality, which would then generate low levels of education, and so on. The society would be trapped in a high-inequality/low-education situation — an 'inequality trap'⁷⁵ In contrast, an equal society would feature the reverse type of equilibrium, equally stable: low inequality would lead to broad educational access, which would lead to low-skills premia, which would in turn ensure that inequality remains low and education levels high.

An inequality trap of the sort just described leaves extensive room for policy intervention, both for normative and efficiency reasons. From a normative

75 See Bourguignon et al. (2007).

point of view, some people are stuck as there is no equality of opportunity: life opportunities and life trajectories are strongly affected by one's family background. People from poor backgrounds are more likely to remain poor, just because they have a poor family. Moreover, the source of wealth differences is questionable from a normative point of view: the rich end up being rich by virtue of being scarce, whereas the poor end up being very poor due to their over-abundance, which in turn is inherited from past inequalities. It seems difficult to justify large differences in welfare on such a basis.

From an efficiency point of view, inequality traps also feature serious problems. First and most obviously, education levels are low compared to what they could be, as testified by the 'good' equilibrium with low inequality and high education levels. This, in turn, directly leads to less productivity, possibly less innovation and possibly large and unnecessary unemployment levels for the unskilled. A second, more subtle, efficiency cost of the inequality trap emerges from the type of people that do and do not obtain education. In an inequality trap, the criterion for obtaining education is family wealth, not ability. The poor are unable to fund their education, even if they have high ability. At the same time, returns to education are very high because of the skills scarcity. Thus, rich families find it worthwhile to educate their children, even if these children are of low ability. This can entail an important efficiency cost: the inequality trap implies that people who do not really benefit much from education (the low ability rich) go to college, while a large proportion of those that would actually benefit most (the high-ability poor) do not. This, again, contrasts with the low-inequality 'good' equilibrium, where the skill premium is low, and the poor earn sufficiently high wages to afford education. In this case, there is no payoff for rich families with low ability children to send them to college, as the benefits are small. Thus, those that go to college are the high-ability individuals, for whom education pays off the most, be they rich or poor.

In an inequality trap situation, both education subsidies and income redistribution can lead to long-lasting social improvements. Both types of policy, by improving educational access below the top of the distribution, can generate a cascade effect which can enable the society to escape the inequality trap. Increased levels of education result in skills being less scarce, and wage differentials become less pronounced. This reduces inequality and improves the capacity of the poor to further increase their access to education, which in turn leads to more skills, less inequality and ever more access, until the 'good equilibrium' is reached.

An alternative policy option, which emerges from this type of framework, is to enlarge the offer of intermediate skills. In a world with only two types of skill, say, highly skilled and unskilled, there is scope for inequality traps implying very large

wage differentials.⁷⁶ The reason is that in a situation with very high inequality, the poor will be so poor that they will not be able to take advantage of the returns to education, even if these are enormous. In contrast, with numerous skill levels, these very large wage differentials are difficult to maintain. Consider an inequality trap with just two skill levels — a strong skill scarcity and large wage differentials between the highly skilled and the unskilled. Now suppose that intermediate semi-skilled training opportunities, with more affordable tuition costs, emerge. Because the semi-skilled families will be richer than the unskilled, they will certainly want to take advantage of the high-skill premium, unless the returns to semi-skilled work are also high. But if this is the case, then the poor will want to exploit the opportunities of semi-skilled work, which they can access more easily due to their lower tuition costs. In any case, skill levels would tend to rise, leading in turn to lower wage differentials, that is, lower inequality.

Inequality and investment with social externalities

Absent from the previous frameworks is the idea that individuals are affected by their social environment in a manner relevant to their investment decisions. Considering these social effects increases the potential for ~~inequality traps, with low investment,~~ to occur. The reason is that individuals are affected by their environment but they also affect their environment in turn. Thus a situation may arise where a ‘bad environment’ leads to low individual investment, which in turn feeds back into the ‘bad environment’. These types of situation lead to specific policy implications, which differ from the ones discussed earlier.⁷⁷

Following Durlauf (2006), social externalities can be categorised on the basis of the type of social environment that one considers. These range from peer groups where individuals actively interact, to reference groups with which individuals identify, to neighbourhoods which might contain both peer and reference groups and where relevant collective decisions, such as the level of school fees, might be taken.

Peer groups, reference groups and neighbourhoods can have a significant effect on an individual’s education decisions. Consider first peer groups. Childhood peers often serve as source of information and inspiration. In this way, they can affect substantially one’s perceptions of returns to education, as well as one’s aspirations. Moreover, to the extent that individuals within a group engage in

⁷⁶ See Mokherjee and Ray (2003).

⁷⁷ The theoretical literature on social externalities and investment include Benabou (1993); Benabou (1996); Fernandez and Rogerson (1996); Durlauf (1996); Brock and Durlauf (2001). See Durlauf (2006) for an excellent non-technical summary of this literature.

imitative behaviour, or to the extent that the group rewards and punishes certain types of behaviour, one's actual behaviour will be influenced by the behaviour of the group as a whole. This can severely affect the development of abilities. For instance, certain peer groups may encourage hard work while others condemn it, thus creating incentives for members of the group to behave accordingly.

Reference groups can have similar effects to peer groups. Even with limited interaction with members of one's reference group, this type of group can have a great influence as a role model. Role models partly shape one's values and aspirations; for instance, role models that practise and condone behaviours which are not conducive to productive investment are likely to discourage educational investments.

Finally, neighbourhood composition can be relevant for individual education decisions in a variety of ways. First, they partly determine peer groups and reference groups, thus helping to shape perceptions and aspirations. Second, they represent the setting in which important types of collective action are taken. One instance of collective action at the neighbourhood level, with important implications for education, is school quality. Residents of certain types of neighbourhoods may be able or willing to invest more heavily in education quality (for instance by setting higher fees) than those in other neighbourhoods. In this way, neighbourhood characteristics exert an externality on the development of each child, regardless of the specific situation of that child. Bright children that would derive strong benefits from quality education might not be able to find schools that match their needs if they live in a neighbourhood in which high quality education is not available.

These types of social externalities can naturally give rise to inequality traps. Certain groups may end up reproducing behaviour that is not conducive to educational investment, while others generate opposite outcomes because of the feedback mechanisms between the group as a whole and each individual's behaviour. Peer groups that encourage ambition, and have optimistic perceptions about the returns from education and productive behaviour, also generate incentives for each of the individuals in the group to exert effort and invest. This behaviour will then tend to confirm the appropriateness of these messages, hence solidifying the outcomes in the group. The same type of mechanisms operate in reference groups, except that the feedback effects tend to occur between different generations, where older generations serve as role models for younger ones, who, in turn, engage in behaviours that lead them to become similar types of role models for the future generation.

Absent from the previous discussion is the issue of group formation. For group externalities to generate inequality traps, groups must themselves be stable; in other words, there must be mechanisms to ensure that individuals stick to their groups, even if these generate bad outcomes. For reference groups, this is achieved

largely by identity and cultural influences which may be difficult to break. Peer groups, in turn, may be strongly constrained by availability, particularly in highly segregated communities. To the extent that neighbourhoods provide the main pool of potential peers, and to the extent that neighbourhoods are themselves stable in terms of characteristics, peer groups will also need to be constrained correspondingly. This brings us to the question of the stability of neighbourhoods themselves and their potential to generate inequality traps; that is, the interaction between residential choices and educational outcomes. Several articles have modelled these interactions.⁷⁸

It emerges from this literature that residential segregation can emerge naturally due to the differential preferences of the rich and poor. The sensible assumption is made that everyone benefits from a 'good' environment (that is, an environment composed of wealthier and more-educated neighbours or of better quality education). Then, to the extent that the rich/educated derive a sufficiently higher benefit from a 'good' environment or that they find it less costly to procure, residential choices lead to segregation. This can arise via rents or via school-financing arrangements. Rich individuals end up living in communities in which rentals and schools are more expensive. Because the rich benefit from a better environment, they are willing to pay more for it and the high rentals and expensive schools deter the poor from moving into those neighbourhoods. In this way, segregation results.

This type of framework delivers specific policy implications. First, it clearly leaves room for policy intervention, since undesirable outcomes at the group level cannot be resolved by individuals acting in isolation. This is because these very individuals act rationally and follow their own interests, given the configuration of the group and the influences it exerts on them. This framework calls for policy interventions at the group level. Only by breaking the potentially harmful group effects can bad outcomes be avoided. This stands in stark contrast to the framework of capital-market imperfections — where redistribution of income by itself could help to improve educational access and break the inequality traps. Here, it is not the lack of income per se that constrains poor individuals, but an environment that depresses their aspirations and their perceptions of the benefits of education. Redistributing income does not help to solve these issues. Instead, policies targeted at neighbourhoods — such as generous public financing of schools in poor communities — can help to break the vicious circle by directly helping the development of abilities, making schooling more attractive to peers and providing better, future role models.

78 See, for instance, Benabou (1996) and Fernandez and Rogerson (1996).

In addition, some of the models in this framework give valuable insights into segregation. On the one hand, they make it clear that segregation may, to a certain extent, be inevitable, short of a continuous enforcement of integration. At the same time, they argue that the amount of segregation that occurs spontaneously will typically be socially undesirable.⁷⁹ This is because individuals, when making their residential decisions, take into account how prospective neighbourhoods will affect them, but not how they themselves will affect the neighbourhoods. Benabou (1996) argues that a highly educated neighbour will typically have a stronger positive impact on a poorer community than they would in one where everyone is already highly educated. Thus, it would be socially desirable to increase the number of highly educated individuals in poorer communities. Fernandez and Rogerson (1996) argue that if the poorest among the rich living in the rich neighbourhood moved to the poor neighbourhood, both neighbourhoods would become richer and hence lead to higher overall investment in education. Thus, a policy that makes poorer neighbourhoods more attractive would be desirable.

Empirical evidence on the relationship between inequality and higher education

Introduction

In this section, we present some brief analyses of our own and complement these with some of the existing findings from the relevant literature in the South African context. We briefly argue that the ‘perfect credit markets’ hypothesis should be rejected in favour of the ‘imperfect credit markets’, using the National Income Dynamics Study (NIDS) data. We then consider whether the social externality models also seem to have validity in South Africa. For the second part, we use Cape Area Panel Study (CAPS) data. The evidence (in either direction), which relates to this second question, is empirically quite weak. Identifying peer effects and social externalities empirically is extremely difficult, and our findings are based on a mixture of theoretical insights, crude summary statistics and plausible conjectures.

In general, our overall analysis involves estimating the differences in various outcomes between groups defined by their relative position in the income distribution. The outcomes in which we are interested relate to the returns to education, educational attainment and career aspirations. Finally, we look at how poorer and richer students fare in terms of scholastic ability, which reflects both neighbourhood and peer effects, as well as other dimensions of school quality.

79 See, specifically, Benabou (1996) and Fernandez and Rogerson (1996).

Data

For our empirical section, we make use of data from Wave 1 of the National Income Dynamics Study (NIDS), conducted in 2008, and Wave 1 of the Cape Area Panel Study (CAPS). The total sample is large, about 30 000 individuals, and the sampling frame is nationally representative. Wave 1 of CAPS involved a cross-section of about 4 800 young adults aged between 14 and 22 years, living in the Cape metropolitan area in 2002. These data are well suited for some of our analyses as they focus in detail on the youth who are at the stage of their lives at which they are finalising their educational decisions. Included in CAPS are several questions about role models, educational expectations and career aspirations. CAPS also includes a standardised literacy and numeracy module, which we used as a combined measure of external school quality and peer effects.

Throughout our analyses, we make use of a simple measure of relative income. Since this is the running variable throughout this section, it is worth explaining how it is calculated and what it does and does not capture. The variable that we used is the quintile in which a respondent's household's per capita income would rank. To explain — suppose that we knew the amount of money which each household earns in aggregate and the number of people in that household. The per capita value is simply the amount of money available per person if they shared the aggregate evenly within the household. We then ranked all households in terms of their per capita income from smallest to largest. We took the poorest 20% of households and placed them in the category called quintile 1. Quintile 2 comprised those households that fell between the 21st and 40th percentiles, and so on, until we reached the richest households between the 81st and 100th percentiles in quintile 5.

What is attractive about this measure is that it cleanly and simply captures groups in terms of their relative economic well-being. This is essentially what inequality is about. What it doesn't do, is reflect the absolute level of well-being corresponding to each quintile. It also doesn't give an idea of how far apart the various quintiles are in terms of levels. However, we know that South Africa was classified as an upper-middle-income country by the World Bank, with a gross national income (GNI) per capita of \$6 090 in 2010. In addition, the World Bank reported a Gini coefficient for South Africa of 63.1 in 2009. This places South Africa as one of the most unequal countries in the world. This means that quintile 5 can be considered to be the 'rich', with developed levels of income, while the remainder can be considered to be the 'poor', with developing country levels of income. Of course, between quintiles 1, 2, 3 and 4 there are varying degrees of poorness and, in some ways, those in quintile 4 seem to have a clearly middle-class lifestyle. Nonetheless, insofar an income matters for our outcomes, we expect

a difference in order of quintile rank, and potentially a much larger difference between quintiles 4 and 5 compared to the other adjacent quintiles.

Empirical findings

How high are returns to a post-secondary qualification in South Africa?

An inequality trap emerging from the imperfect credit-markets hypothesis implies that the returns to education remain exceptionally high due to the scarcity of skilled people. In Table 6.1, we consider the cohort of 25–29-year-olds in NIDS in 2008. We separate them into three educational categories, and estimate their mean employment rates and wages conditional on employment.⁸⁰ Assuming that the wage from being unemployed is zero, one measure of the rates of returns to a college education as compared to a matric only, would be to calculate the products of the employment probability and mean wages for the two groups, and to then calculate the ratio of these products.

Turning to the table, we first observe that the fraction with any type of tertiary qualification in this group is approximately one in six. In fact, the largest group are people that never finished secondary school, at about three in five. This would be expected if unskilled workers could get reasonably well-paid jobs easily, but the employment rates of high-school dropouts is below 50%, while those with some tertiary schooling is 30 percentage points higher at 73.6%. Moreover, the mean income, conditional on employment, goes up dramatically as we increase the educational qualifications. From the least educated to those with only a matric, the mean income more than doubles. From those with only a matric to those with more than a matric, it more than doubles again. Note also that this is for a fairly young group, and these disparities would almost surely get wider with time. This means that our estimate of the returns to education are likely to be conservative.

Finally, if we consider the ratio of expected values, getting a matric will triple your expected earnings, while moving from a matric to some tertiary qualification will increase your expected earnings by about 167%.

To conclude, our analysis does find extremely high returns to education, as well as a relatively small fraction of people obtaining a tertiary education. These findings are widely supported by the findings of several other researchers. For example, Lam (1999) found evidence of increasing returns to incremental levels of education in South Africa. Hoogeveen and Özler (2005) and Leibbrandt Levinsohn and McCrary et al. (2010), at the household and individual levels respectively, both identified increases in the rates of return to education between

80 All summary statistics are calculated including the post-stratification weights released with the surveys.

1995 and 2001. Thus, taken together, the empirical evidence is consistent with the existence of an inequality trap driven by credit constraints.

Table 6.1: Educational attainment and returns to education

25–29-year-olds in NIDS Wave 1					
Education	(1)	(2)	(3)	(4)	(5)
Level	N	% (weighted)	Employed (%)	Total income (mean)	Expected value (3)*(4)
No matric	1 028	59.94	43.2	1 237	534
Matric only	356	23.55	56.9	2 681	1 524
Some tertiary	173	16.51	73.6	5 547	4 081
Total	1 557	100			

Notes:

1. The total income column is calculated using only 656 observations. This is due to:
 - a. The low levels of employment (754 observations in our sample).
 - b. Missing data/invalid responses.
2. All summary statistics are calculated using the post stratification weights.
3. The definition of employment includes wage employment, self employment, casual work, working without pay in a family business, or subsistence agriculture.

Is attainment related to income?

A second testable implication of the credit-constraints hypothesis is that the rich are more likely to attain a tertiary education. In Table 6.2, we present the mean proportion with some completed tertiary qualification, by income quintiles, for each age in the 20–24-year-old age cohort. At age 20 years, the young adults in the richest households are four times more likely to have some tertiary qualification than any of their counterparts. While differences between the lower quintiles are not systematic, there is a clear advantage to being in quintile 4, compared to quintiles 1, 2 and 3. Moreover, despite the gradual catching up with age of the youth in the lower quintiles, by age 24 years, the proportion of youth in the richest households remains double that of those in the second richest quintile, at 31.4%. Again, we find that the evidence is clearly consistent with the credit-constrained model.

Table 6.2: Mean of some tertiary qualification

(By age and income quintile: NIDS Wave 1)						
	Per capita household income quintiles					
Age	1	2	3	4	5	Total
18	0.000	0.000	0.007	0.000	0.000	0.001
19	0.000	0.041	0.004	0.005	0.037	0.016
20	0.050	0.036	0.038	0.054	0.212	0.070
21	0.035	0.106	0.034	0.043	0.290	0.087
22	0.035	0.063	0.026	0.117	0.345	0.101
23	0.005	0.089	0.007	0.130	0.233	0.084
24	0.097	0.123	0.098	0.157	0.314	0.151
Total	0.027	0.058	0.031	0.075	0.191	0.068

Why do youth who do not have a tertiary qualification not enrol?

A third dimension that is central to the theoretical models being considered is that credit markets do not work well. Given the extremely high rates of return from education, even average students would rationally want to get some type of a tertiary qualification. Yet the fraction that does so is fairly low, especially in the bottom three quintiles. In Table 6.3 we summarise the primary reasons given by 18–24-year-olds for not enrolling even though they do not have any tertiary qualification. As usual, we do so separately by income quintile.

The most important categories include having found employment, financial costs and looking for employment, all of which are directly or indirectly related to financial status. However, there is considerable heterogeneity when we consider the responses by quintiles. In the poorest two quintiles, the main reasons are the costs of staying in school and the decision to look for a job, while the fraction of respondents who stopped enrolling because they had actually found a job is relatively low. Fertility-related reasons also feature strongly in this group. The pattern is similar for those in quintile 3, but a substantial fraction had stopped enrolling because they were working instead of looking for work. Quintile 4, our ‘middle-class’ quintile, are once again somewhere between the poorer quintiles and the top quintile. While in the poorer quintiles 1 to 3, somewhere between one in three and one in four respondents left their education due to the financial costs of remaining in school; in quintile 4 this number drops to about one in six. In the top quintile, this fraction decreases even further to one in ten. Moreover, more than half of the ‘dropouts’ from the top quintile do so because they already have a job. This fraction is more than double the overall average.

This fits well with other existing research. In a very comprehensive report by Branson et al. (2009), the authors found that the returns to tertiary education are indeed very high, but that students face two substantial, yet quite different, sets of constraints that restrict their ability to obtain such qualifications. First, many students choose subjects or attain grades at the pre-tertiary level which result in them not being eligible to enter into tertiary institutions. Second, the costs and financial considerations associated with attending a tertiary institution can be prohibitively high, for most of the (minority of) students whose secondary school performance does meet the various entrance criteria.

In summation, we find clear evidence that returns to education are very high, that richer students are much more likely to attain a tertiary qualification, and that financial costs are a significant deterrent to poorer students. Overall, the data seem to strongly favour the credit constraints model over the perfect credit markets model.

Table 6.3: Reasons for not being enrolled in 2008 (NIDS Wave 1)

	Quintiles					Total
	1	2	3	4	5	
Finished school/education	9.1	15.1	12.2	10.5	16.1	12.1
I was working	3.7	7.4	21.7	35.4	52.2	21.3
Could not afford to stay in school	32.2	24.8	31.2	17.4	10.0	24.7
Wanted to look for a job	24.8	27.1	14.7	18.3	11.6	19.9
Was pregnant/had a baby	17.2	13.5	9.1	9.2	3.8	11.3
Was needed at home	2.0	1.6	2.8	1.1	0.3	1.7
Was ill/sick	2.7	2.5	2.8	2.1	0.6	2.3
I got married	2.1	2.3	0.0	1.7	0.7	1.4
Grades were very poor	2.1	0.4	2.4	1.5	0.6	1.5
Was suspended/expelled	0.9	0.5	0.3	0.5	1.1	0.6
Other (specify)	3.3	4.7	2.3	1.8	3.2	3.0
Too old	0.0	0.0	0.5	0.5	0.0	0.2
Total	100	100	100	100	100	100
N	422	396	440	350	135	1743

Notes:

1. Sample is aged 18–24 years, are not currently enrolled and do not have a tertiary qualification
2. Means are weighted using the post-stratification weights

What empirical evidence is there to support the social externalities theories?

In general, it is very hard to find compelling evidence that speaks to the social externalities hypotheses. The reason for this is that, due to social stratification by neighbourhoods, it is almost impossible to disentangle neighbourhood effects from individual or household-level income effects. Even if one were to focus on the small number of poorer students who commute to richer schools, any comparison would be contaminated by unobserved selection; the parents who make such sacrifices are likely to have a strong preference to invest in their children's education. This would probably manifest in several other dimensions which are important but unobservable to an empirical economist.

To date, there are at least two papers that speak to potential peer effects. The first, by Lam et al. (2009), makes use of CAPS data and finds that the age profile of one's peer group significantly affects the probability of an early sexual debut for girls. The second, by Garlick (2012), makes use of the random allocation of peers in dormitories at the University of Cape Town to investigate the effects of the scholastic aptitude of one's peers on one's own academic performance at university. He finds that peer effects are substantial and significant, and manifest most strongly for weaker students.

We provide two pieces of evidence that speak to the social externalities literature, one relating to career aspirations and the other concerning the development of abilities.

In Table 6.4 we summarise CAPS data and calculate the fraction responding to a particular category as their response to the question, 'What work do you expect/plan to do at age 30?' Recall that these were all young adults aged 14–22 years at the time of the Wave 1 survey in 2002. The responses are varied as there are several categories, so we present the largest categories and collapse the remainder, which is the majority of responses, in an 'other' category.

Overall, it is hard to see clear patterns in this table. We find, for example, that youth in the poorest quintiles were the most likely to plan to be doctors, lawyers, social workers, college professors and nurses. Youth in the 5th quintile, in contrast, were most likely to be in agriculture or fisheries, decorators or designers or general managers in trade. The lack of a clear pattern, combined with the fact that even the poorest youth expect to be doctors and lawyers, suggests that aspirations are not being negatively affected by social externalities.⁸¹

81 In analyses of the NIDS data that are not included in this chapter, we also found no discernible differences in terms of expectations with regard to obtaining a tertiary qualification across the quintiles.

Table 6.4: Expected occupation at age 30 years. (CAPS Wave 1 data)

Occupation	Quintile					Total
	1	2	3	4	5	
Accountant	4.68	3.64	4.4	7.44	6.3	5.39
Medical doctors	7.75	6.47	4.85	2.93	4.14	5.14
Agriculture and fishery	4.29	4.12	6.08	3.15	7.63	5.11
Lawyers/attorneys	6.56	4.86	5.43	3.88	2.89	4.61
Social worker	7.27	4.81	3.16	3	0.8	3.67
Teaching professional in higher education	4.95	4.47	1.68	2.63	2.25	3.18
Nursing	6.26	4.27	3.78	1.87	0.53	3.17
Decorators and designers	1.39	1.57	1.81	3.23	6.69	3.14
Machinery, mechanics and fitters	2.43	3.88	2.78	2.65	1.5	2.6
Mechanical engineers	2.48	3.35	2.61	2.74	1.21	2.42
Productions and operations manager	2.5	2.13	2.68	2.74	1.79	2.34
General manager (wholesale/ retail trade)	1.5	1.67	1.55	1.86	3.29	2.04
Other	47.94	54.76	59.19	61.88	60.98	57.19
Total	100	100	100	100	100	100

The second piece of data relates to scholastic aptitude, and how it varies across income groups. We use the standardised literacy and numeracy scores from CAPS for this, and present the results graphically. In order to make the graphs clearer, we collapsed the respondents into three income categories: ‘low income’ which means that they are in quintiles 1 or 2; ‘middle income’ which corresponds to quintiles 3 or 4; and ‘upper income’ which corresponds to quintile 5. We then plot the distribution of aptitudes for three different age groups. This allows us to compare the differences in aptitudes across income groups, as well as how these differences evolve with time.⁸²

The first thing that is striking is that the distributions are clearly different, and unambiguously improve with income. Even among the youngest CAPS respondents, these differences are already pronounced by the time they would ordinarily be in the early parts of secondary school. This corroborates the findings of Branson et al. (2009) regarding the statement that many students might

82 Technically, we cannot separate between the effects due to aging or potential differences in cohorts.

experience sufficiently poor learning environments well before the tertiary level, such that eligibility for the tertiary level becomes a binding constraint.

When we consider the distributions among slightly older respondents, the differences become even more pronounced. The distributions for the low- and middle-income groups are mostly stable, with slight increases in both the means and variances. In contrast, the distribution of the upper-income youths clearly shifts sharply to the right, and converges to the upper bound of the test.⁸³

These findings are likely to be a result of several factors, one of which must be school quality. School quality itself is likely to be a function of resources, infrastructure and teacher quality, as well as peer effects. As explained earlier, we cannot separate the effects of these factors empirically. A partial resolution might be found in Wittenberg (2005). He makes use of time-use data from Statistics South Africa's Time Use Survey 2000 and analyses differences in time allocated to school and studies by children from different socio-economic groups. He finds that punctuality and absenteeism seem to be disproportionate problems among poor learners, and that poor learners spend considerable time each day on chores. One could interpret this as weak evidence in favour of the social-effects hypotheses, although the links are somewhat tenuous.⁸⁴

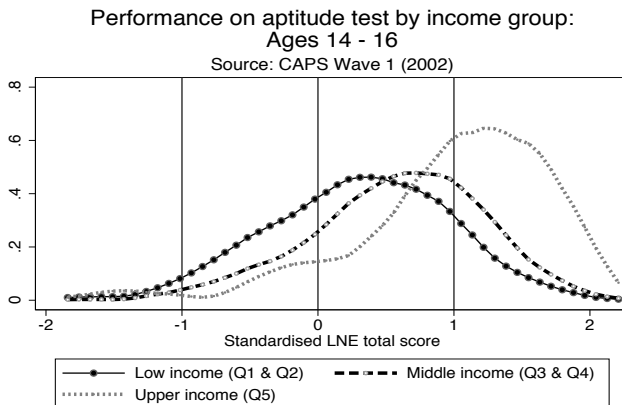


Figure 6.1: Performance on aptitude test by income group, ages 14–16 years
Source: CAPS Wave 1 (2002)

83 The test was relatively simple and comprised just a few questions. For this reason, it is a relatively poor discriminator of aptitude among relatively strong students. We would expect that a more thorough test would have yielded even more striking levels of divergence as the groups from different socio-economic backgrounds got older.

84 It could also simply be the case that chores in richer households are outsourced to domestic workers, or are less time-intensive. This would be a more conventional resources-based explanation.




Figure 6.2: Performance on aptitude test by income group, ages 17–19 years
Source: CAPS Wave 1 (2002)




Figure 6.3: Performance on aptitude test by income group, ages 20–22 years
Source: CAPS Wave 1 (2002)

Policy discussion

There appears to be sufficient evidence to suspect that South Africa is caught in an inequality trap where high inequality leads to low levels of skill accumulation, which in turn consolidates the high levels of inequality. The trap works particularly through tertiary education: it is at the tertiary level that access is very limited and that returns are very large. Credit constraints and social externalities are both likely to play a role in sustaining this trap, although we do not find evidence that lack of aspirations are part of the story. Credit constraints and social externalities matter for access to tertiary education directly, when making the actual decision, as well as indirectly, by limiting access to high-quality education at the primary

and secondary levels. This in turn prevents students from qualifying for tertiary studies. Thus, policies that seek to break the South African inequality trap need to address access to tertiary education, as well as the large differences in education quality at the primary and secondary levels.

We consider first policies that focus directly on improving access to tertiary education. One branch of these policies needs to address the issue of credit constraints. Theoretical models that incorporate credit constraints often mention general progressive redistribution as a policy implication. In the context of South Africa, this message should probably be viewed more as an argument adding to the debate on the benefits and costs of redistribution, than as a practical means to actually increase access to education. The reason is that in South Africa, credit constraints in the tertiary decision are likely to apply to a minority of the population and are unlikely to bring about dramatic improvements in access: most prospective students from poorer backgrounds are actually not eligible to enter into universities. To this minority, however, credit constraints appear to be binding, and targeted redistributive programmes are likely to be effective. Currently, there is an extensive means-targeted public programme of financial aid, the National Student Financial Aid Scheme (NSFAS), which provides bursaries and loans. In principle, this should alleviate financial constraints for students willing and eligible to attend higher education.

However, even in the presence of this programme, there are reasons why these constraints may still bind. First, prospective students often appear to lack information regarding the NSFAS. Second, even if eligible for public financial aid, students typically need to pay registration fees up front. The existence of these fees deter applications to higher education institutions. Thus, for the public programme to be more effective, efforts should probably be directed towards extending the financing to cover registration fees, as well as disseminating information regarding the benefits and procedures of the programme more widely.

A second policy intervention, which flows from our analysis, is to strengthen intermediate degrees, that is, qualifications between the high school certificate and a university degree. As mentioned earlier, enlarging the offer of intermediate skills can help to reduce inequality and promote skill acquisition. In South Africa, technical, vocational education and training (TVET) colleges in principle fulfil this role. However, there are reasons to believe that access to these colleges could be improved. ~~The focus groups mentioned earlier~~ provide some evidence that a lack of information and the issue of registration fees might be most problematic with regard to TVET colleges. First, discussions made clear that prospective students have very little information on the types of studies and the financing possibilities afforded by public TVET colleges. As an example, several students had applied only to the TVET colleges that had visited their school, which often happened to

be private colleges which typically charge higher fees. Second, there is evidence that school performance is a poor indicator of skill in many South African schools (see Lam et al. 2011). This suggests that students face high uncertainty regarding the level of their own skills and the likelihood of being accepted into universities. In this case, registration fees might deter applicants, particularly to TVET colleges. When constrained by registration fees, students might put their resources into university applications rather than TVET colleges in order to avoid the frustrating eventuality of finding themselves eligible for university but having not applied.

Alleviating credit constraints and encouraging access to TVET colleges would help to improve skills and reduce inequality. However, to achieve a major improvement, eligibility constraints need to be addressed. This involves encouraging broad access to quality education at the primary and secondary level. There are two basic approaches to this. First, to increase integration of the poor into existing high-quality schools, and, second, to improve the quality of poorer schools.

We consider first the challenge of integration of students from poorer backgrounds into high-quality schools. As mentioned earlier, the same type of constraints that prevent the poor from accessing higher education — that is, credit constraints and social externalities — can deter access to high-quality schools. High-quality schools demand not only higher transport costs and fees, but also higher expenses to ensure successful social integration, (for example, via extra-curricular activities and lifestyle choices). Peer pressures (from one's own as well as other social groups) and role models may make it difficult for students from poor neighbourhoods to integrate successfully into schools dominated by students from rich backgrounds. Policy interventions to alleviate these constraints could include fee waiving for students from poor backgrounds (a policy already in place), as well as additional aid to cover transport and extra-curricular activities.

There are reasons to believe that these types of policies may achieve only limited success. First, it is difficult to conceive and implement policies that limit negative social influences on students from poorer backgrounds. Second, as argued by Selod and Zenou (2003), rich (typically white) families can adapt their behaviour to counter the effects of such policies. In their model, white parents raise fees to counter policy interventions to encourage black attendance to high-quality, formerly white schools. More generally, privileged families may increase all sorts of monetary and non-monetary barriers to outsiders. They may even leave the system altogether to set up private schools, where the potential of government intervention is limited, taking with them the very resources (for instance high-quality teachers) which made their schools of higher quality in the first place. Moreover, integration policies, even if successful, are likely to drain poorer neighbourhoods of their most able students, thus deepening potentially

negative social externalities in the schools of those neighbourhoods. Finally, and possibly most importantly, this type of policy can succeed in improving the skills of only a minority of poor students. The reason is that high-quality schools in richer neighbourhoods are so few in relative terms that they can productively absorb only a small proportion of students from poorer backgrounds.

A policy geared towards increasing the quality of poorer schools has more potential to be successful. Based on the evidence provided here, as well as the work of several others, it seems widely accepted that improving school quality remains a fundamental developmental objective, which is also frustratingly difficult to achieve. The post-1994 era involved a shift in investment and redistribution towards historically disadvantaged schools, but research by Van der Berg and Burger (2003) suggests that this did not result in improved student performance. However, for a combination of methodological and interpretative reasons, we should interpret their findings with caution. Indeed, Case and Deaton (1999) found quite the opposite—that school resources in South Africa do, indeed, explain differences in scholastic performance.

One aspect of improvement, which we feel may be very important for improving school quality, is to improve teacher quality. When we looked at career aspirations in CAPS, we found that fewer than 10 respondents, out of approximately 4 800, plan to be teachers at the age of 30. Teaching is clearly not seen as an occupation of choice for most young people. There are several reasons why this might be the case, and each of these could be addressed as part of a holistic approach to improving teacher quality. First, to qualify as a teacher generally requires a four-year university degree. As argued in this chapter, this comes with considerable costs and barriers. Second, the financial remuneration is not that attractive, given the cost of the qualification. Starting salaries for teachers were below R95 000 per annum in 2015, and while this could go up with seniority and qualifications, they remain relatively low when compared to private-sector employment which also requires a four-year degree. As such, the financial calculus and the time investment, combined with the stream of financial payoffs, may explain the low ranking of teaching as a career of choice.

More generally, there are other mechanisms that may also improve teacher quality. First, it might be worthwhile to reward teachers who display high levels of initiative and motivation. Second, there should be some negative consequence for poor behaviour. In the gift-exchange models of effort (Akerlof, 1982; 1984), people decide on what is fair effort for a given wage, and their concept of fairness is calculated by some reference group. Even if a small fraction of teachers are putting in low effort and obtaining the same wage, this could make other teachers resentful, and eventually the equilibrium might converge to a low-effort equilibrium.

There may be political constraints to implementing overall wage increases or incentive schemes for teachers. Some social actors traditionally oppose the former, while other actors oppose the latter. A possible politically feasible package would be to combine the two policies as a bundle: higher overall wages and incentive schemes. This could become palatable to a wide spectrum of social actors while promoting overall teacher quality.⁸⁵

Moreover, our arguments, particularly those relating to remuneration, link almost directly back to inequality — there is a scarcity of high-quality education. Skilled people have to choose between medium-paid, skilled occupations (which have positive social externalities) or highly paid, skilled occupations in the private sector (which do not have the same social externalities), and this subsequently guarantees the lack of human capital in the following generation. This vicious circle is stable or possibly even widening, and will continue to replicate itself unless strong and well-directed government intervention is forthcoming.

References

- Aghion, P. & Bolton, P. (1997). 'A theory of trickle-down growth and development', *Review of Economic Studies*, 64 (2): 151–172.
- Akerlof, G. (1982). 'Labour contracts as partial gift exchange', *Quarterly Journal of Economics*, 97 (4): 543–569.
- Akerlof, G. (1984). 'Gift exchange and efficiency-wage theory: Four views', *American Economic Review*, 74 (2): 79–83.
- Banerjee, A. & Newman, A. (1993). 'Occupational choice and the process of development', *Journal of Political Economy*, 101: 274–298.
- Benabou, R. (1993). 'Workings of a city: Location, education and production', *The Quarterly Journal of Economics*, 108 (3): 619–652.
- Benabou, R. (1996). 'Equity and efficiency in human capital investment: The local connection', *Review of Economic Studies*, 63 (2): 237–264.
- Bertola, G. (2000). 'Macroeconomics of distribution and growth'. In: A.B. Atkinson & F. Bourguignon (eds). *Handbook of Income Distribution*, 1st edition, volume 1. New York: Elsevier, 477–540.
- Bertola, G., Foellmi, R. & Zweimüller, J. (2006). *Income Distribution in Macroeconomic Models*. Princeton, NJ: Princeton University Press.
- Bourguignon, F., Ferreira, F. & Walton, M. (2007). 'Equity, efficiency and inequality traps: A research agenda', *Journal of Economic Inequality*, 5: 235–256.
- Branson, N., Leibbrandt, M. & Zuze, T.L. (2009). *The Demand for Tertiary Education in South Africa*. SALDRU Research Report, University of Cape Town.

85 Note that these suggestions, while directed specifically at teachers, would also apply to any other professions that generate strong social externalities. Examples of these professions would include nurses, other health-care professionals and social workers.

- Brock, W. & Durlauf, S. (2001). 'Discrete choice with social interactions', *Review of Economic Studies*, 68 (2): 235–260.
- Case, A. & Deaton, A. (1999). 'School inputs and educational outcomes in South Africa', *The Quarterly Journal of Economics*, 3: 1047–1084.
- Durlauf, S.N. (1996). 'A theory of persistent income inequality', *Journal of Economic Growth*, 1 (1): 75–93.
- Durlauf, S.N. (2006). 'Groups, social influences and inequality: a memberships theory perspective on poverty traps'. In: S. Bowles, S.N. Durlauf & K. Hoff (eds). *Poverty Traps*. Princeton, NJ: Princeton University Press.
- Fernandez, R. & Rogerson, R. (1996). 'Income distribution, communities, and the quality of public education', *The Quarterly Journal of Economics*, 111 (1): 135–164.
- Foellmi, R. & Zweimüller, J. (2006). 'Income distribution and demand-induced innovations', *Review of Economic Studies*, 73 (4): 941–960.
- Galor, O. & Zeira, J. (1993). 'Income distribution and macroeconomics', *Review of Economic Studies*, 60: 35–52.
- Garlick, R. (2012). *Academic Peer Effects with Different Group Assignment Rules: Tracking versus Random Assignment*. Available from: http://www-personal.umich.edu/~rgarlick/garlick_peereffects_tracking.pdf (accessed 24 June 2012).
- Ghatak, M. & Jiang, N. (2002). 'A simple model of inequality, occupational choice, and development', *Journal of Development Economics*, 69: 205–226.
- Giuseppe, B., Reto, F. & Josef, Z. (2006). *Income Distribution in Macroeconomic Models*. Princeton, NJ: Princeton University Press.
- Hoogveen, J.G.M. & Özler, B. (2006). *Not Separate, Not Equal: Poverty and Inequality in Post-apartheid South Africa*. William Davidson Institute Working Paper No. 739, University of Michigan.
- Lam, D. (1999). *Generating Extreme Inequality: Schooling, Earnings, and Intergenerational Transmission of Human Capital in South Africa and Brazil*. PSC Research Report No. 99-439, August, Population Studies Center, University of Michigan.
- Lam, D., Ardington, C. & Leibbrandt, M. (2011). 'Schooling as a lottery: Racial differences in school advancement in urban South Africa', *Journal of Development Economics*, 95 (2): 121–136.
- Lam, D., Marteleto, L. & Ranchhod, V. (2009). *Schooling and Sexual Behavior in South Africa: The Role of Peer Effects*. PSC Research Report No. 09-694, December, Population Studies Center, University of Michigan.
- Leibbrandt, M., Levinsohn, J. & McCrary, J. (2010). 'Incomes in South Africa after the fall of apartheid', *Journal of Globalization and Development*, 1 (1): Article 2.
- Leibbrandt, M., Woolard, I., Finn, A. & Argent, J. (2010). *Trends in South African Income Distribution and Poverty since the Fall of Apartheid*. OECD Social, Employment and Migration Working Paper No. 101. Paris: Directorate for Employment, Labour and Social Affairs, OECD.
- Maoz, Y.D. & Moav, O. (1999). 'Intergenerational mobility and the process of development', *Economic Journal*, 109: 677–697.
- Mookherjee, D. & Ray, D. (2003). 'Persistent inequality', *Review of Economic Studies*, 70: 369–393.

- Owen, A.L. & Weil, D.N. (1998). 'Intergenerational earnings mobility, inequality and growth', *Journal of Monetary Economics*, 41: 71–104.
- Piketty, T. (1997). 'The dynamics of the wealth distribution and the interest rate with credit rationing', *Review of Economic Studies*, 64: 173–189.
- Piketty, T. (2000). 'Theories of persistent inequality and intergenerational mobility'. In: A.B. Atkinson & F. Bourguignon (eds). *Handbook of Income Distribution*, 1st edition, volume 1. New York: Elsevier, 429–476.
- Selod, H. & Zenou, Y. (2003). 'Private versus public schools in post-apartheid South African cities: Theory and policy implications', *Journal of Development Economics*, 71 (2): 351–394.
- Van der Berg, S. & Burger, R. (2003). 'Education and socio-economic differentials: A study of school performance in the Western Cape', *South African Journal of Economics*, 71 (3): 496–522.
- Wittenberg, M. (2005). *The School Day in South Africa*. CSSR Working Paper No. 113, Centre for Social Science Research, University of Cape Town.
- World Bank. (2006). *World Development Report: Equity and Development*. Washington, DC: World Bank.