The Dual Role of Income in the Spread of HIV in Africa

Chris Desmond
2002

The Dual Role of Income in the Spread of HIV in Africa

Chris Desmond
2002

In fulfilment of a Masters of Commerce
October 2002

University of Natal

Supervisor
Professor D. Posel

Declaration

I, Chris Desmond, declare that this thesis represents original work that has not been previously submitted in any form to any university. Where use has been made of the work of others, this has been duly acknowledged and referenced in the text.

Signed:_____

Date:_____

Acknowledgments

I would like to thank UNICEF South Africa for funding the initial research. I would also like to thank Prof. Cornia of UNICEF IRC Florence and Dr Patel of UNICEF East and Southern Africa Kenya for their useful and insightful comments on early drafts.

Thanks also go to my family for their help and to Emelda for her patience.

A special thank you and acknowledgment to my supervisor Dori Posel for great help and support throughout the preparation of this thesis.

Abstract

Every day more people in Africa are infected with HIV despite prevention efforts. These new infections and those already infected are not evenly spread throughout the continent. Substantial variations in HIV prevalence exist within and between countries. Understanding these variations helps understand what is driving the epidemic and this understanding in turn helps in the design of more appropriate interventions to prevent its further spread. This thesis builds on existing work by attempting to develop a more comprehensive theory of what role income plays in the spread of HIV. To this end the Theory of the Dual Role of Income in the Spread of HIV is outlined and explained. It uses the concepts of relative and absolute income, borrowed from elsewhere in the health economics literature, to separate the different effects income has on individual and group risk of HIV infection. The theory hypothesises that, while higher levels of absolute income (income independent of others) offer protection against infection via better access to health care and information, higher relative income (the income of an individual relative to other members of their social or reference group) increases risk of infection either as a result of more sexual partners or higher risk partners. The theory in no way argues that HIV infections are not related to poverty, but rather that the relationship is somewhat more complicated and non-linear than often suggested. The explanatory power of the theory is examined with the use primarily of two data sets: firstly using data collected from antenatal clinics in two South African provinces linked with census data and secondly with data on a large South African company. While these data are not ideal, the results from the analysis are in line with the expectations based on the theory. The theory and the results of the analysis presented in this thesis support the argument that environments in which decisions are made and actions taken are important in determining risk of HIV infection. This argument suggests that prevention efforts need to do more then provide information.

Table of Contents

Declaration ii

Acknowledgements iii

Abstract iv

Table of Contents v

List of Figures viii

List of Tables ix

Introduction 1

Chapter 1 Literature Review 11

1.1 Introduction 11

1.2 Explaining the differential spread of HIV 13

1.2.2 Factors influencing the probability of transmission 15

1.2.3 Factors influencing the number of contacts 17

1.2.4 Individual behaviour and risk 23

1.2.4.1 Sexual behaviour studies 23

1.2.4.2 Who is HIV positive? 25

1.2.5 Factors influencing group prevalence 27

1.2.6 Explaining the contradiction 28

1.3 Income and the health economics literature beyond AIDS 30

1.4 Summary 33

Chapter 2 The Theory of the Dual Role of Income in the Spread of HIV 35

2.1 Introduction 35

2.2 The role of absolute income in the spread of HIV 38

2.2.1 Individuals' absolute income 38

2.2.2 Average group absolute income 43

2.2.3 The strength of the relationship 46

2.2.4 The onset of the epidemic 50

2.2.5 Summary of the role of absolute income 52

2.3 Relative income 53
2.3.1 Male relative income 57
2.3.2 Female relative income 58
2.4 Combination 62
2.5 Observed relationship between income and HIV risk 66
2.7 Summary and hypotheses 68

Chapter 3 Method and data 74
3.1 Introduction 74
3.2 Review of existing studies 74
3.3 An ideal data set 75
3.4 Antenatal clinic data combined with the 1996 census 80
3.4.1 Data 80
3.4.2 Analysis 87
3.4.2.1 The education variable 91
3.4.3 Summary of antenatal clinic data analysis 92
3.5 Company X data 93
3.5.1 Data 93
3.5.2 Analysis 97
3.5.3 Summary of company X data analysis 98
3.6 Additional company data 100
3.7 Summary 100

Chapter 4 Results 103
4.1 Introduction 103
4.2 Re-examination of existing studies 104
4.2.1 Individual risk of infection 104
4.2.2 Country studies 108
4.2.3 Summary 109
4.3 Antenatal clinic data 109
4.3.1 The age variable 112

4.3.2 Female status variable 113
4.3.3 Male short distance work related mobility 114
4.3.4 Monthly household expenditure 116
4.3.5 Education 117
4.3.6 The model 119
4.3.7 Summary 120
4.4 Analysis of Company X data 121
4.4.1 Age variables 124
4.4.2 Provincial site location 125
4.4.3 Ethnic group 125
4.4.4 Accommodation 126
4.4.5 STIs 127
4.4.6 Income 127
4.4.7 The model 130
4.4.8 Summary 131
4.5 Analysis of further company data 132
4.6 Summary 133

Chapter 5 The Extension of the Theory of the Dual Role of Income in the Spread of HIV 136
5.1 Introduction 136
5.2 Individual changes in income 136
5.3 Income distribution 138
5.4 Sub-group changes in income 140
5.5 Summary 141

Chapter 6 Conclusions 143

References 147

Appendix 154

List of Figures

0.1	HIV prevalence among women attending antenatal clinics in South Africa – 2001	2
0.2	The uneven spread of HIV in KwaZulu-Natal	3
0.3	HIV epidemic- determinants	5
1.1a	HIV incidence in a simulated generalised epidemic	18
1.1b	HIV prevalence	18
2.1	HIV risk and access to health services	48
2.2	The relationship between absolute income and the risk of HIV infection	49
2.3	Absolute income and risk at the onset of an epidemic	51
2.4	Male risk and relative income	58
2.5	Male and female risk and relative income	61
2.6	Simple combination of absolute and relative income	63
2.7	Decreasing role of relative income	64
2.8	Absolute and relative income with overlapping reference groups	65
2.9	Predicted observable relationship between HIV prevalence and income	67
2.10a	Absolute and relative income and individual risk of infection	70
2.10b	Predicted observable relationship between HIV prevalence and income	71
3.1	Shortcomings of the linear probability model	90
4.1	Male probability of dying	105
4.2	Female Probability of dying	106
4.3	Probability of infection and years of education	118
4.4	Income and predicted risk of infection	128
4.5	The effect of income on probability of infection where STI prevalence and accommodation vary	129
4.6	HIV prevalence by job band	132

List of Tables

3.1	Antenatal clinic data descriptive statistics	86
3.2	Company data descriptive statistics	96
4.1a	Logit regression results assuming the entire sample as the reference group	110
4.1b	Logit regression results assuming the local community as the reference group	111
4.2	The probability variation in risk of infection for a move from the minimum to maximum value of each variable	120
4.3	The results of the logistic regression on company data	122
A1	Variables examined but excluded from the final model	153

Introduction

While the world awaited the coming of a new millenium on 31 December 1999, over 34 million of its inhabitants did so already infected with HIV (UNAIDS, 2000). Almost 19 million more had already died of AIDS, three million of them during that year, and over 13 million children spent the night without their mother because she was among the dead.

Prevention campaigns continue, promoting abstinence from sex, faithfulness to one partner and the use of condoms. Despite these efforts, UNAIDS estimates that over five million people were infected with HIV in 1999 alone. The focus of prevention activities suggests that changing individual choices through the provision of information is the key to turning the tide of this epidemic. The findings contained in this thesis contribute to the growing evidence that it is rather the environment in which decisions are made, and in which behaviours occur, that is driving the HIV epidemic. There are cultural, social and economic factors that shape behaviour and the risk of that behaviour. Without understanding and addressing these factors, the provision of information alone will have limited success.

It is not an objective of this thesis to develop a prevention intervention; its aim is rather to contribute to the understanding of the socio-economic factors driving the epidemic, with a focus on what role income plays in the spread of HIV. Through this contribution, it is hoped that this research, along with other work of a similar nature, will help inform future efforts to prevent the further spread of HIV.

The importance of socio-economic factors in the spread of HIV is highlighted in the examination of the uneven spread of the disease. UNAIDS estimate that 24 million of the 34 million people living with HIV at the end of 1999 lived in Africa, and the prevalence within the adult population of countries of sub-Saharan Africa varied from 1.35 per cent in Niger to 36 per cent in Botswana (2000). Even within countries, HIV prevalence can

vary greatly. In South Africa, provincial HIV prevalence among women attending clinics in 2001 varied from less than 10 per cent in the Western Cape to over 33 per cent in KwaZulu-Natal (Department of Health, 2002). Figure 1 displays the levels of HIV infection in the nine South African Provinces.

Figure 0.1 HIV prevalence among women attending antenatal clinics in South Africa – 2001

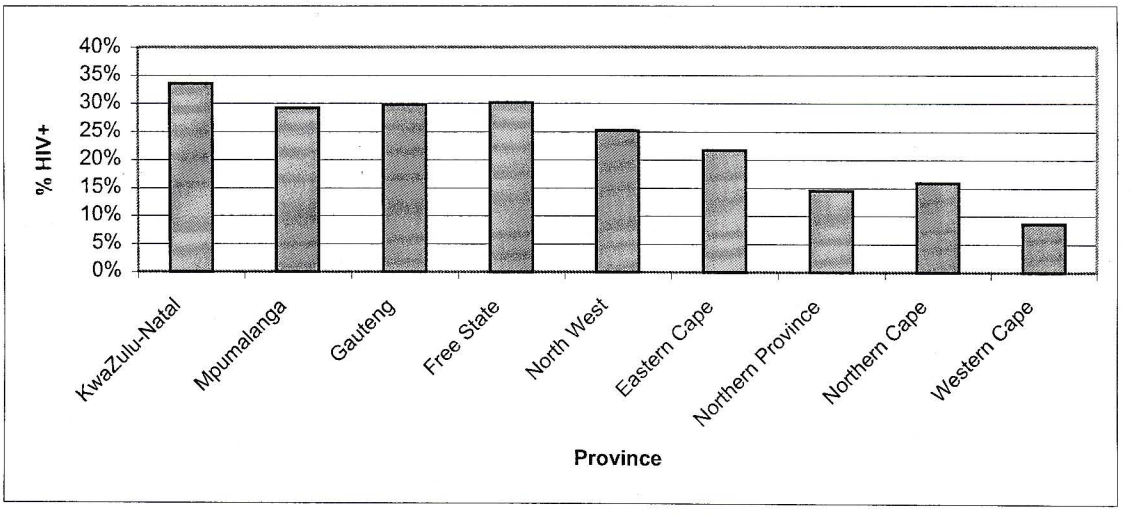

Infection rates vary widely at an even more micro level. In 1999 in KwaZulu-Natal, one of South Africa's provinces, for example, HIV prevalence at antenatal clinics varied from 10 per cent to close to 50 per cent. Figure 2 highlights the variation between the different health districts within the province.

Figure 0.2 The uneven spread of HIV in KwaZulu-Natal

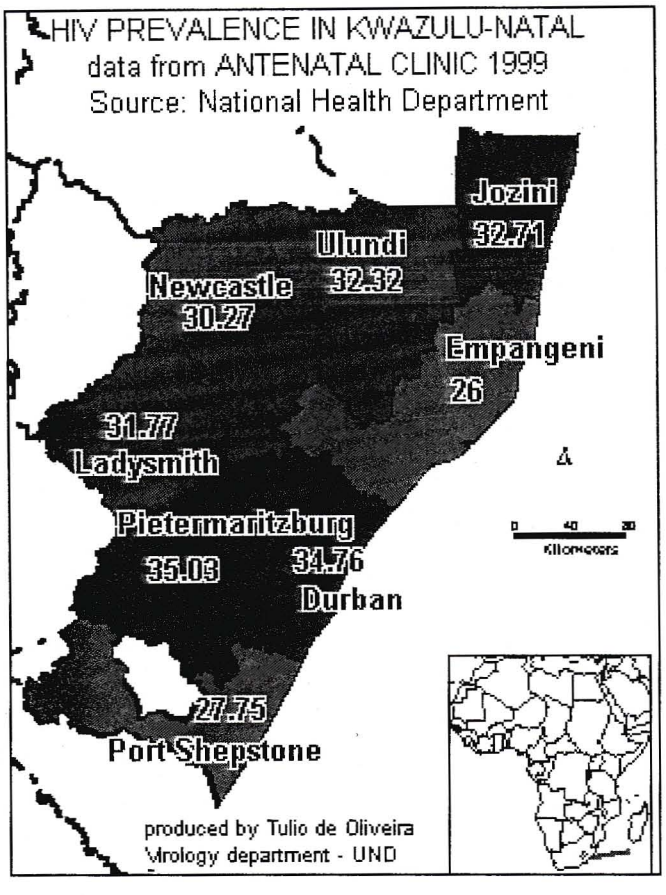

These dramatic variations in the spread on HIV, between and within continents and countries, have prompted researchers to investigate the causes. The resultant research has suggested a range of possible factors leading to the variations. One of the major causes for the variation between continents is that there appears to be a number of different epidemics, with different characteristics and at different levels of maturity, unfolding around the world. HIV in Africa is largely a heterosexual epidemic with many men and women infected and new infections occurring primarily through unprotected vaginal sex. HIV infection rates have reached far higher levels than anywhere else in the world. By the end of 1999 there were 16 countries in Africa in which HIV prevalence was in excess of 10 per cent of adults aged 15-49; in seven of those prevalence was close to, or in

excess of, 20 per cent (UNAIDS, 2000). In Asia, by contrast, there are only three countries where adult prevalence exceeds 1 per cent.

In Latin America, HIV has three major pathways of transmission (UNAIDS, 2000). While there is a growing heterosexual epidemic, infections levels are also high among men who have sex with men and among intravenous drug users. In the high income regions of the world, the bulk of infections are within the male homosexual population and the drug using population (UNAIDS, 2000).

Given this variety in the nature and maturity of the various HIV epidemics in the world it would have been impossible to address the role of income across all of them within one thesis. The role of income is likely to be very different for drug-user epidemics as opposed to those driven by sexual transmission. Furthermore power relationships between the sexes associated with income obviously lead to different roles for income in heterosexual epidemics compared to homosexual ones. While there may be some cross cutting issues it was necessary to select one type of epidemic to focus on. The literature reviewed and the theory developed in this thesis, therefore, focus on the epidemic in Africa. The African epidemic was selected because of its maturity as one of the issues discussed in this thesis relates to how the role of income may change as the epidemic matures.

In Africa, differences in migration patterns and levels, poverty rates, STI prevalence, and the incidence of male circumcision and condom use are some examples of factors identified in the literature as influencing the spread of the disease. Chapter one of this thesis discusses, in detail, this type of research, the methods applied and the different determinate of variation identified. For now, however, the following figure provides a simplified framework for understanding some of the important factors.

Figure 0.3 HIV epidemic – determinants

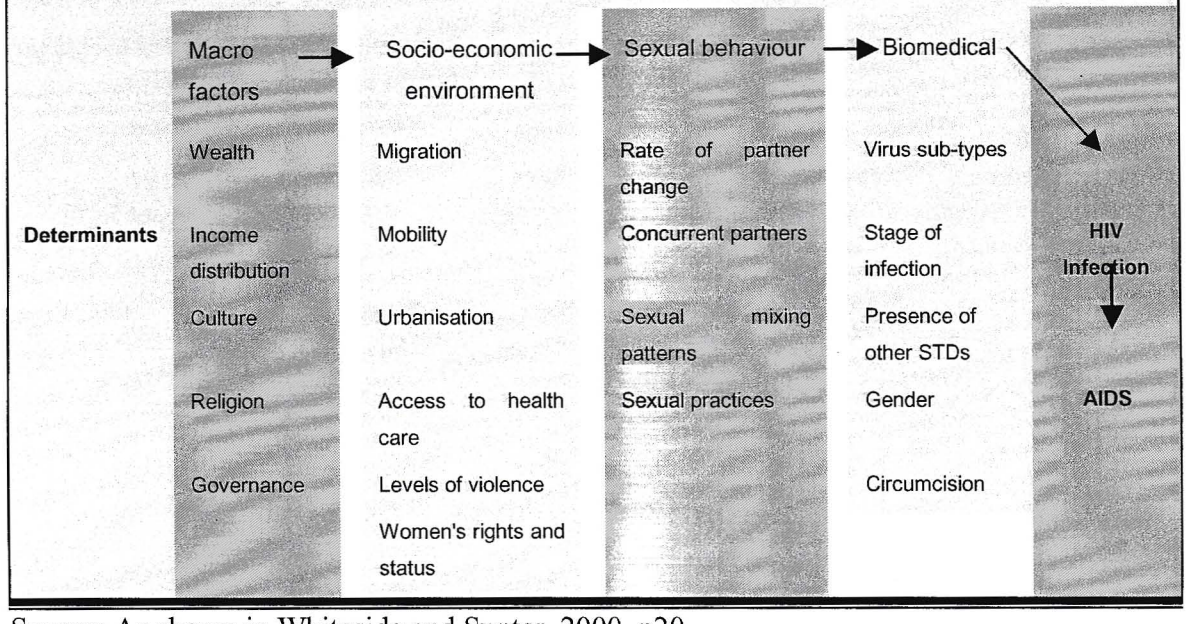

Source: As shown in Whiteside and Sunter, 2000, p20.

HIV is transmitted in Africa primarily through heterosexual intercourse. The risk of each sexual act is determined by a variety of biomedical factors, which in turn are the result of sexual behaviour. Such behaviour, however, occurs within the context of the broader socio-economic environment, which is itself shaped by the macro environment. The above is a simplification, as the relationships are not so clearly ordered. Indeed, the macro environment may affect biomedical factors directly; for example, in a Muslim country the macro environment might lead to higher rates of male circumcision and therefore an affect on a biomedical risk factor. Regardless of how different influences affect risk, the point remains that, although sexual in nature, HIV infections occur in a broader socio-economic framework.

For many of the factors associated with HIV infection the relationship is not unidirectional. HIV infections and the resultant AIDS deaths themselves have an impact on socio-economic conditions and other variables that affect the risk of HIV infection. The relationship between HIV infection and some of these factors appear to work in both directions. This is certainly the case regarding income. This thesis examines the link between income and HIV infection, from the point of view of how income affects risk of

infection. There is, however, also an impact on income and socio-economic status as a result of HIV infection. Whiteside (2002) discusses how, once infected, an individual faces increasing costs associated with their illness at the same time as their ability to work is reduced. Such an impact is felt not only by the individual, but by the household in which they live, and collectively HIV infections are, and will increasingly be, affecting entire communities and countries. While the impact of HIV infections on income, and indeed on many other variables, is an extremely important area of research, this thesis concentrates on the relationship between income and HIV infection up to the point where infection has occurred.

Although simplified, figure 0.3 already highlights that there are many important factors contributing to the spread of HIV and accounting for the uneven spread. While it is important to develop an understanding of as many of these as possible, that is well beyond the scope of this thesis. This thesis focuses on one of the contributory factors, the role of income in the spread of HIV.

A considerable amount of insightful and useful work on the relationship between HIV and income has already been done internationally, notably Ainsworth et al, 1998, Cohen, 1993 and 1998, Stillwaggon, 2000 and Over, 1998. This thesis builds on this work in an effort to develop a more in-depth understanding of the role of income in the spread of HIV. It is argued that this greater depth of understanding can be obtained by applying concepts borrowed from elsewhere in the health economics literature. With the use of the concepts of relative and absolute income, commonly used in debates on differential health outcomes between developed countries, the work already done can be developed into a more comprehensive set of explanations. Absolute income refers to an individual's, or a group's, income independent of other individuals or groups. Relative income is the income of an individual relative to other members of their group. An individual may be relatively wealthy although in absolute terms they are poor. This can occur if their group is poor but they have income above the average of their group. While still being poor, they are considered to be relatively wealthy when compared to other members of their group.

The basis of the more comprehensive explanation is that income has a dual role in relation to risk of HIV infection. Absolute income, it is argued, offers protection against infection through, among other things, better medical treatment for STIs and better access to information. Relative income, however, is argued to affect access to resources and power within a group; therefore members of the group with higher relative income have greater ability to sustain multiple partnerships, at least in the case of men, and as a result incur greater risk. It is argued that the average absolute income of a group of individuals determines the average risk of the group, but that an individual's income relative to other members of that group determines their risk relative to that average. Furthermore, the two types of income are argued to work in opposite directions: the higher the group's average income in absolute terms, the lower their average risk, but the higher an individual's income relative to others within the group the higher their risk relative to the average. Such reasoning suggests that the relatively wealthy of poor groups will have the highest infection rates.

The above explanation is obviously an oversimplification as there are many complications, such as who should be considered to be in a group together and how the effects of income do, or do not differ for men and women. The following chapters attempt to deal with some of these complications as the concepts of relative and absolute income are applied to the question of the relationship between HIV and income. The results of this application in this thesis are then presented as the theory of the dual role of income in the spread of HIV. The predictions of this theory are then examined in relation to the analyses of a number of data sets. To this end the structure of this thesis is as follows.

Chapter one: Literature Review

This chapter provides reviews of two sets of literature: firstly, studies relating to factors affecting the spread of HIV in Africa and, secondly, literature discussing the concepts of relative and absolute income. The research presented in this thesis has a similar focus as

that found in the first set of literature and the theory is based largely on the findings of previous work in this area. The chapter attempts, with the use of a simple model, to place the different studies and the factors they identify in relation to each other. This process highlights the complex interplay between different variables in determining group and individual risk. While other research is mentioned, the review focuses on studies that have examined the links between income and socio-economic status and risks of HIV infection.

The second section of the review discusses the literature on the concepts of relative and absolute income. These concepts and the theories based on them have been developed primarily in efforts to explain why average health status differs between wealthy countries with similar incomes. While there is no intention to enter into such a debate in this thesis, a review of the literature helps in gaining an understanding of the concepts that are borrowed from it.

Chapter two: The Theory of the Dual Role of Income in the Spread of HIV

Drawing on the two sets of literature reviewed in chapter one, this chapter presents an attempt to build on the literature on the socio-economic determinants of risk by borrowing the concepts of relative and absolute income and thus developing a more comprehensive explanation of the role of income in the spread of HIV.

The chapter presents the Theory of the Dual Role of Income in the Spread of HIV, which argues that the effects of relative and absolute income on the risk of HIV infection work in opposite directions: higher levels of absolute income decrease the average risk of groups, but individuals within groups with higher levels of relative income, that is those at the upper end of the income distribution within a group, have greater risk.

Chapter three: Method and Data

This chapter outlines the methods used and describes the data to which they were applied in the process of examining the validity of the theory developed in the preceding chapter. The analysis consists of four sections, each of which is outlined and discussed in this chapter. The first section involves the comparison of the results of two previous studies reviewed in chapter one, with expectations based on the theory developed in chapter two. The second section involves regression analysis using individual level data from two provinces in South Africa from the yearly antenatal clinic survey of HIV prevalence. The chapter discusses how the antenatal clinic data set was combined with data from the 1996 South African Census to examine the role of relative income. The third section discusses how data collected on HIV prevalence within a large South African company was used to investigate the empirical evidence of the protective effect of absolute income. Finally, the chapter discusses the use of data from other companies in the region and highlights the limitations faced in drawing conclusions from it.

Chapter four: Results

The results chapter outlines the findings of the analyses described in the previous chapter. As with the proceeding chapter, the results are divided into four sections: firstly, the comparisons between expectations based on the theory developed in chapter two with the findings of two existing studies; secondly, the results of the regression analysis using antenatal clinic and census data; thirdly, the results of the regressions conducted using the company data and, finally, the discussion based on the results from other companies.

For each section of analysis the results are presented, discussed and compared with the hypotheses generated in chapter two. While the data did not allow for the examination of all of the hypotheses generated in chapter 2, they provided support for those that were examinable. The findings in all four sections were in line with expectations.

Chapter five: Extension of the Theory of the Dual Role of Income in the Spread of HIV

In chapter two the concepts of relative and absolute income were applied to the problem of explaining what role income has in determining the risk of HIV infection. The theory developed in the chapter is, however, static in so far as it attempts to explain the role income has played in determining the current pattern of risk. It is, however, possible to extend the theory and apply the framework developed in chapter two to a dynamic analysis. This chapter presents such an extension of the theory. While the results from chapter four are in line with some of the expectations based on this extension, appropriate data to examine the quality of predictions based on the extended theory do not exist.

Chapter six: Summary and conclusions

The final chapter of this thesis summarises the process and results and draws out conclusions. It is argued in this chapter that this thesis provides a more detailed and comprehensive analysis of the role of income in the spread of HIV. The separating of the relative and absolute effects of income allows for a more flexible analysis, which was able to build on the work previously done in the area.

Furthermore, the chapter argues that this thesis adds to the growing evidence that causes of the uneven spread of HIV are both contextual and behavioural and, even when behavioural, they are behaviours that must be understood in relation to the context in which they occur and which shapes them.

Chapter 1

Literature Review

1.1 Introduction

The variation in the spread of HIV between and within countries and continents has prompted many researchers and academics to examine the factors that have lead to such differences. Related to this type of investigation is the research that attempts to determine who within a group is most at risk of infection. Research and examinations in these areas have brought to light a number of issues and informed policies for preventing the further spread of HIV.

All the factors identified influence the spread of HIV in two possible ways. An individual's risk of becoming infected is determined by the number of times they come into contact with the virus multiplied by the risk of transmission per contact. Some factors increase the number of contacts, while others increase the risk per contact, and some may influence both.

This review examines and discusses the literature that has developed in an effort to identify the factors that influence the spread of HIV. The discussion is shaped by continually returning to the basic formula of individual risk i.e. risk = contacts * risk per contact.

The purpose of the review is two fold: firstly, to place this thesis in the context of the literature and to highlight the problem that this investigation seeks to address; secondly, to develop an understanding of concepts used elsewhere in the health economics literature that will be applied in later chapters to the question at hand. This thesis examines in detail, the role of income in the spread of HIV, investigating how risk associated with different levels of income may differ between and within reference

groups. The first part of this chapter will outline the existing arguments and analysis relating to income and HIV spread and will compare the results and conclusions. What is of interest is that the studies addressing *individual* risk find that as income increases so does the probability of HIV infection (see for example Ainsworth et al, 1998; Allen et al, 1992; Kirunga et al, 1997). At a *country* level, however, studies have found that poorer states have higher infection rates (Cohen, 1998). Informative efforts have been made to reconcile these differences (World Bank, 1998). The discussions presented in this thesis, however, argue that through the application of concepts developed in relation to differential health levels in the developed world a more in-depth reconciliation and understanding can be achieved.

This thesis will draw on current debates and theories in the health economics literature relating to the role of income in the determination of health at group and individual level, in an attempt to expand on existing research on the role of income in the spread of HIV and the apparent contradiction between individual and country studies. While the first part of this chapter discusses literature relating specifically to HIV, the second part outlines and discusses the relevant health economics literature relating to the importance and role of income in the determination of health generally. This is not intended as a review of the findings, but rather of the concepts that have been developed in the literature, which will be applied specifically to the role of income in the spread of HIV in the following chapters. Research dealing with the link between health and income has been undertaken for many years (Deaton, 1999). Recently, however, considerable debate has taken place in journals and research papers around the development of a more detailed understanding of how income influences health and whether income distribution is as, if not more, important than income itself (Deaton, 2001a). This debate, although not directly relevant to the problem at hand, provides a number of useful concepts that can be used to better understand the role of income in the spread of HIV. These concepts include relative and absolute income, reference groups, and neighbourhood impact on health.

1.2 Explaining the differential spread of HIV

There are a number of different HIV epidemics in the world: the characteristics of the mainly homosexual Western Europe and North American epidemics are very different from the primarily heterosexual epidemic in Africa, although the epidemic among poorer Americans bears some resemblance. HIV in Western Europe and North America has been spread primarily through homosexual sex and intravenous drug use (UNAIDS, 2000). The situation is similar in Eastern Europe and the former Soviet countries (UNAIDS, 2000).

It is not possible in this research to examine the role of income in the spread of HIV in all contexts, as the situations and possible responses are very different. This thesis focuses on the African heterosexual epidemic. The majority of HIV infections are in Africa and the need to understand the epidemic here is essential and urgent, not only for the continent, but for others looking to avoid replication. Although there are some similarities, many of the factors influencing the spread of HIV in other continents are likely to be very different from those in Africa. Furthermore, the epidemic in Africa is far more advanced and the levels of infection far higher (UNAIDS, 2000) The role of different factors may change at different stages of the epidemic and investigating new and advanced epidemics simultaneously would complicate an already difficult issue. The review of the broader literature concerning other factors influencing risk and spread of HIV is also concentrated on studies conducted in Africa. Unless otherwise stated, the literature in the rest of this section is concerned with the African epidemic.

This thesis examines the role of income in the spread of HIV. This role, however, cannot be considered in isolation and must be placed in the context of the literature examining the social and biomedical factors contributing to, or inhibiting, the spread of HIV. The factors relating to income are linked to the many other issues identified in the literature discussed here, and an understanding of these issues is essential for the later analysis. Understanding how the different factors influence one another and indeed how the literature itself fits together is difficult and at times requires background knowledge of

the epidemic and the terminology used. To facilitate understanding, this section includes some discussions on the necessary knowledge to link the literature together and provide the background to the subsequent chapters.

The literature discussed here can be understood on the basis of the following equation and function:

Prob(Individual HIV) = contacts * probability -(1)

Prev(Group) = f(Average (Prob(Individual HIV))) -(2)

Where:

Prob(Individual HIV) is the probability that an individual will become HIV positive

Contacts is the number of times an individual comes into contact with HIV

Probability is the probability of infection per contact

Prev(Group) is the HIV prevalence of the group

HIV can be transmitted in a variety of ways. In the African epidemic, characterised by heterosexual transmission, equation (1) can be interpreted as the number of sexual contacts an individual has with HIV positive individuals, multiplied by the probability that the virus will be transmitted during each act. It is these two factors on which the literature concentrates: factors that increase the number of times individuals come into contact sexually with an HIV positive person and those that influence the probability of transmission per sex act.

This review discusses the social science literature. The large body of medical literature on this topic is beyond the scope of this study. Where medical information is relevant, it is considered in the context of social science research. For example, circumcision has been identified in the medical literature as a factor that reduces the probability of HIV transmission per sexual act. In this review, this is discussed in terms of studies that have examined the prevalence differences between communities with high rates of circumcision and those with low rates, this as opposed to medical research outlining why circumcision reduces the probability of infection. Similarly, studies are discussed

regarding sexually transmitted diseases and dry sex (the practice of using drying agents in the vagina), again examining the outcome of dry sex for the spread of HIV rather than the medical/physical reasons for the increase in risk.

The factors influencing the frequency of sexual contacts with HIV positive persons and the probability per contact are roughly divided in the literature into bio-medical factors and behavioural considerations. Bio-medical factors essentially address the probability of infection per act, and behavioural factors the number of contacts. Behavioural issues, however, may also affect the probability of infection per act; anal sex for example is a behaviour that increases the risk of infection per sexual act.

1.2.2 Factors influencing the probability of transmission

Studies in industrialised countries have estimated the probability of HIV transmission per sexual act to be small, although greater for women. The average probability that HIV will be transmitted from an infected female to an uninfected male during unprotected vaginal sex, is approximately 1 in 1000, and 2 in 1000 from an infected male to an uninfected female, according to medical studies (Haverkos et al, 1992). Anal sex has a higher risk, as a result of greater friction and bleeding, with between 5 and 30 infections per 1000 acts (DeGruttola et al 1989). These estimates are, however, likely to be under-estimates as they were based on studies of discordant couples, where one was infected and the other was not; couples where the virus was transmitted quickly were excluded (Mastro et al,1996). Furthermore, the estimates are in the absence of co-factors that increase the risk of transmission, such as STDs. The average probability in contexts where STDs are common are, therefore, far higher (World Bank, 1997).

A number of factors have been identified that increase these probabilities. The presence of one or more Sexually Transmitted Infections (STIs) greatly increases the probability of infection per act, particularly if one or both of the sexual partners has an ulcerative infection, as these facilitate the entry of the HIV through the skin. (For a review of findings, see UNAIDS/WHO, 2000). Numerous studies, both medical and social science,

have identified the role of STIs in the transmission of HIV (see for example chapters 2, 6 and 7 in World Bank, 1998 for a summary of findings and UNIADS/WHO, 2000 for a review of studies).

Dry sex, the practice of using drying agents such as detergents, carbonated drinks and traditional medicines to tighten the vagina in an effort to increase the pleasure of the male partner, has also been linked to an increase in the probability of transmission. Williams et al (2000), in a study in South Africa, found that women who practised dry sex had higher HIV prevalence than those who did not use drying agents. This was true for sex workers (91 per cent compared to 66 per cent) and non-sex workers (50 per cent compared to 41 per cent). In the sample, however, only 3.2 per cent of women used drying agents. The correct use of condoms, as is well known, reduces the probability of HIV infection per act. Male circumcision is also associated with a reduction in transmission probability (Moses et al, 1990). In a recently published study, Auvert et al examined the relationship between HIV infection and male circumcision in four African cities (2001). In one city in Kenya HIV prevalence among circumcised men was 9.9 per cent compared to 26.6 per cent among uncircumcised men. The authors used logit regressions to control for socio-demographic characteristics, sexual behaviour and the presence of other STIs. Having done so, they found that circumcision continues to be associated with lower risk of infection with an odds ratio of 0.26 (1.12-0.56 95 per cent confidence interval).

The factors, such as the presence of an STI, condom use and dry sex, which affect the probability of transmission, are themselves influenced by socio-economic and cultural variables. The factors identified above are the direct influence through which socio-economic and cultural variables affect an individual's risk of HIV infection. Obtaining a better understanding of the direct influences, outlined above, allows for more meaningful interpretations of the studies that address the variables that influence both the number of contacts and the probability per contact. The following section examines the variables that influence the number of contacts an individual has with HIV.

1.2.3 Factors influencing the number of contacts, sexual behaviour and their interplay with the probability of transmission.

If there were a group of adults, all of whom where HIV negative, living on a desert island and each day they all selected a new sexual partner there would still be no HIV, because if there was no HIV to begin with there would be no HIV at the end. The greatest influence on the number of contacts with HIV, and therefore on the prevalence, is the prevalence at t-1 – that is the prevalence in the time period before the current period. HIV follows a typical epidemiological pattern. In the early stages, when HIV has just entered a group[1], new infections over a given time period (the incidence) are few, as few people come into contact with the virus. As more people become infected, however, the probability that an individual will come into contact with an already infected individual increases and the incidence of new infections increases. As the epidemic matures, the growth in incidence slows and eventually falls as the number of new infections per time period falls. This occurs primarily because many of the people who are at risk of infection are already infected. They may come into contact with another HIV infected individual, but no new infections will result. The probability of an HIV negative individual coming into sexual contact with an HIV positive individual will continue to increase even if the incidence of new infections within the group falls because, although new infections are falling, prevalence will continue to rise. Figures 1a and 1b display the relationship between prevalence and incidence in a simulated generalised epidemic.

[1] HIV can enter a group when a member of the group gets infected by a member of another group and then infects others within their own group. This is typically associated with migration, into the group or out and

Figure 1.1a **HIV incidence in a simulated generalised epidemic**

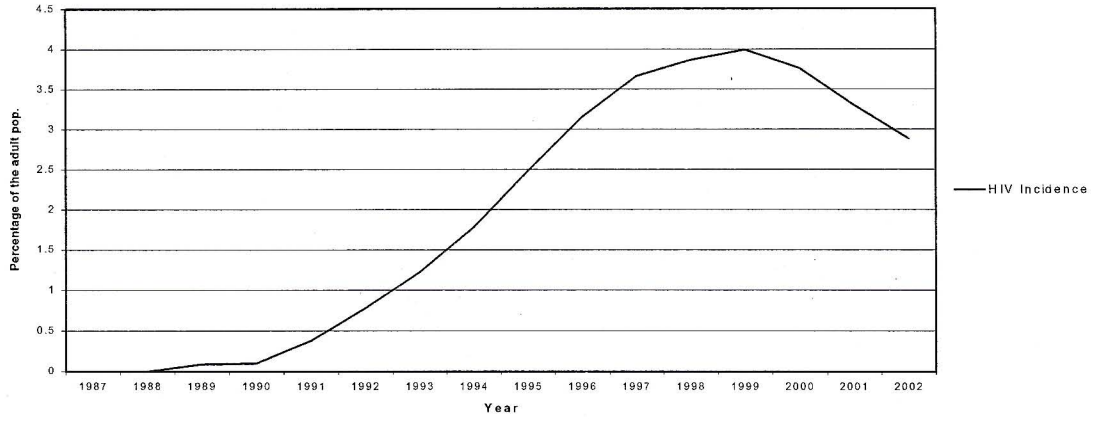

Figure 1.1b **HIV prevalence**

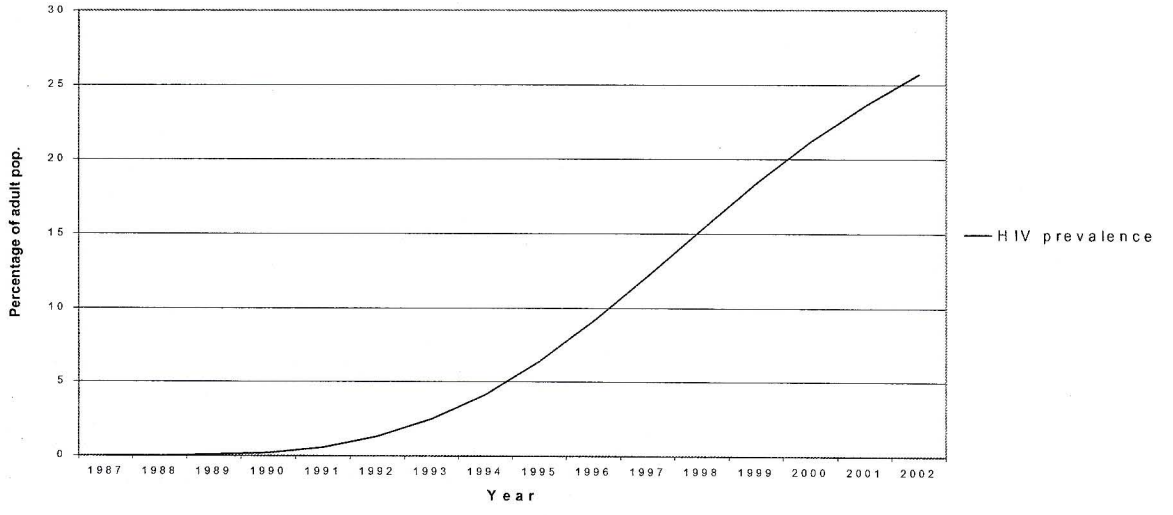

So long as the incidence exceeds the death rate prevalence will rise. As prevalence rises the probability that an HIV negative individual will come into contact with an HIV positive individual increases. The risk of infection for HIV negative individuals, therefore, continues to increase even as the number of new infections decreases.

---

back again.

Prevalence in all time periods is a result of a number of other factors, including those that influence the probability that individuals will come into contact with an HIV-infected person.

At t_0, for an HIV epidemic to begin, the virus must be introduced into a group, community or country. As the virus is transmitted only through the exchange of bodily fluids, there are a limited number of ways in which the virus can be introduced. It could enter a population via contaminated blood products from an infected population, or through sexual contacts between a migrant from that population with another population already suffering from the epidemic or sexual contacts with migrants from populations where the epidemic has begun. Migration is, therefore, clearly an important influence affecting the epidemic (see for example Campbell, 1997, Zwi et al, 1991, and Hunt, 1989 or more recently Lurie, 2000a).

Migrants are important not only in introducing the virus, but also in facilitating its spread. Migrants live away from home, often being away from their families for long periods of time. This separation from home and family is associated with high-risk sexual behaviour and HIV infection (Campbell, 1997). Once infected, migrants then infect partners at home on their return, although recent evidence suggests that this goes both ways (Lurie et al, 2000). Research in rural KwaZulu-Natal and two migrant destinations, involving the testing, for HIV, of couples, one of whom was a migrant, found that in 40 per cent of discordant migrant couples (where only one partner is infected) it was the women at home who were infected (Lurie et al, 2000). It appeared, therefore, that it was, in some cases, the woman who was infected first. This was, however, a study of a mature epidemic, where the communities between which the migrants travelled were all suffering with high levels of HIV infection.

Once introduced, HIV is spread not only by migrants, but by members of groups and communities. The sexual behaviour of individuals in their choice of partners, rate of partner change, concurrent partnerships, for example, greatly affect their likelihood of coming into contact with an HIV-infected sexual partner.

It is the literature on the factors that influence sexual behaviour, and the context in which it occurs, in terms of both the number of contacts and the risk per contact that is relevant to this work, because it is at this stage that the role of income becomes apparent. The following literature deals not only with behaviour and its influence on the number of contacts, but also on the probability of transmission. As the factors under consideration become more distant from the sexual acts that result in transmission, so it becomes more difficult to separate those influencing the number of contacts from those affecting the probability of infection per contact.

Cohen (1998) provides a useful paper highlighting a number of the more important factors. He begins by arguing that, even after many years, there is still very little understanding of the importance of different factors in explaining the variations in the level of HIV prevalence and its distribution in Africa. Having noted our limited knowledge, he discusses some of the factors that have been identified. Poverty is identified as a major contributor to the spread of HIV, a point he made in detail in an earlier paper (Cohen, 1993). Poverty affects HIV transmission via three channels: behaviour, education and health. People living in poverty are constrained in their decision making and may be forced into survival strategies that increase their risk of infection. The most obvious of these are labour migration and commercial sex work, both of which increase the number of sexual contacts and the probability that they will come into contact with an infected partner. In addition to constrained decision making, the low levels of education and access to information, which characterise poor communities, mean that educational information on ways in which to protect oneself will be difficult to access and understand. Finally, the poor often have little access to health services and generally have poor health status as a result of inferior housing, lack of clean water, and poor nutrition. This situation leads to greater risk and length of illness, including STIs, which increase the risk of transmission per contact.

Stillwaggon (2000), in her paper examining the HIV epidemics in Africa and Latin America, goes further to argue that there has been too great a concentration on

behavioural risk factors associated with poverty and that they are not the only issues. She argues that it is also the poor health status of the poor that has facilitated the spread of HIV. She discusses the correlations between calorie intake, protein consumption, and other factors typically associated with susceptibility to infection. On the basis of current medical studies she can only hypothesise and provide supporting evidence from studies of other illnesses, as there is as yet no medical proof of the link which she asserts. Stillwaggon provides further support for her arguments with the use of data collected on Latin American countries. She presents the results of an ordinary least squares estimation with the dependent variable as AIDS cases per 100 000 of the population and independent variables as the real per capita GDP, level of urbanisation, international migration and calorie supply. The analysis involved data from 20 countries. The coefficient on calorie intake was negative, which Stillwaggon argues supports her assertions, but she does admit that it is also a proxy for poverty.

Cohen (1998), however, notes that the relationship between poverty and HIV is not as simple as it may first appear. He points out that many non-poor people are also infected. These people, he argues, are at risk because their higher levels of income lead to higher levels of mobility and greater access to sex.

Structural adjustment programmes and the associated increase in unemployment and decrease in health care and education expenditure are also linked by Cohen to the spread of HIV. He further argues that the development of the Southern African region has for too long been reliant on migration and the disruption of family life. These movements have widened sexual networks, facilitating the spread of HIV between communities.

Finally, he points out the importance of the position of women in society for the spread of HIV. Low levels of access to education, health care and credit, lack of legal protection of property and little political power leave women in a vulnerable position. Economic gaps between men and women appear to lead to sexual exploitation and the greater spread of HIV.

Rivers et al (1999), in their paper examining what it is about gender relations that leads to increased risk of infection, have expanded on the risk associated with inequalities between the sexes and the implications for the spread of HIV. They also provide a good summary of much of the earlier work in this area with a concentration on Africa, but including studies conducted elsewhere. Where there are unequal power relations between men and women it is the decisions of the men that are implemented. In the sexual environment this means that decisions about safe sex would be made by men and that these decisions would be governed by the dominant constructs of masculinity in the society, often promoting multiple partners.

Rivers and Aggleton (1999) argue that, in developing countries, it is the economic dependence on men and women's lack of education that facilitate the growth of unequal power relations. Women are in positions where they are unable to control the where and when of sex (International Centre for Research on Women, 1996). In such situations, women's desire for safer sex may be met with no co-operation by their male partners (Meursing et al, 1995). The World Health Organisation (1994) has pointed to the large literature on the difficulties for women in trying to persuade men to use condoms.

Rivers and Aggleton (1999) go further to discuss what is it that shapes the decision making of male partners. They discuss the importance of existing constructs of masculinity and how men who do not conform to these are viewed as unsuccessful, leading to possible 'cultural exile'. Prevailing understandings of masculinity often encourage behaviour that promotes the spread of HIV, a prime example being multiple partners. These prevailing views influence the sexual health of men and their partners. An example of discrimination against non-conformist men is that experienced by gay men in a number of societies.

Just a cursory review of the conceptual literature shows that there are a number of different factors interacting with one another in contributing to the risk environment. A number of studies have attempted to quantify the relative importance of these factors at both an individual and group level. The following section reviews some of the more

recent work. It begins with studies of individual risk and behaviour and goes on to those looking at the determinants of group risk. The emphasis is on studies examining the socio-economic influences on the risk of infection.

1.2.4 Individual behaviour and risk

The studies in this area can be further divided into two groups: those that examine the determinants of sexual behaviour and those that estimate the risk of infection. The two are linked, in as far as sexual behaviour influences risk and thus the determinants of sexual behaviour influence risk. The two differ, however, in the formers attempt to isolate particular factors. For example, studies on sexual behaviour may attempt to ascertain the affect of a variable, say education, on condom use and on multiple partner numbers, whereas the latter would examine only the net risk associated with education. If the variable has a different influence on condom use than on multiple partners, the studies will have quite different results. Behavioural studies will examine both, while risk studies will identify only the net effect.

1.2.4.1 Sexual behaviour studies

Studies of this nature are fairly rare for the general population, although a number of related studies exist with regard to the commercial sex work industry (see Ahlburg and Jensen, 1998). Filmer (1998) examined the socio-economic correlates of sexual behaviours associated with risk of HIV infection. This, as already mentioned, is different from examining the correlates to HIV infection. Filmer used the most recent data (1991-94) available from the Demographic and Health Surveys (DHS) of nine countries, seven countries in Africa and two in Latin America. Filmer builds on earlier work by Carael (1995) who examined data collected by WHO on 15 different countries. He argues that the Carael study was limited, in so far as much of the analysis was bivariant and dependent on endogenous variables such as marital status. Filmer attempts to improve on this by using unconditional data from a different source to examine the relationships between sexual behaviour and socio-economic characteristics.

Filmer examined the data on recent non-regular sexual partners, condom use with non-regular partners, and whether 15-19 year olds had ever had sex. These variables were examined in relation to age, urban residence, occupation, household assets, community attitudes, and household characteristics. The results showed that men were more likely than women to have a non-regular sexual partner; the difference was even more marked for the probability of having multiple non-regular partners. Filmer also found that in many of the samples the probability of a man having a non-regular partner increased with education. Carael (1995) argues that this is true for both men and women.

Increases in education, however, were strongly related to condom use and decreased risk of HIV. This is an example of where studies of sexual behaviour differ from those of risk. Studies of risk would only measure the net affect of increasing education and not identify the different components - i.e. the increased risk resulting from more partners and the reduction in risk stemming from greater condom use. Filmer notes that the DHS are not designed primarily for the collection of sexual behaviour data and that this limits the analysis. No efforts were made in the collection of data to identify the level of mis-reporting, a problem that is often associated with sexual behaviour studies (Gersovitz et al, 1995). It must also be noted that any cross country analysis must be conducted carefully as the samples, particularly of men, are drawn differently in the different countries.

Deheneffe et al (1998) conducted a similar analysis of earlier data collected from 12 sites/countries by the World Health Organisation between 1989 and 1991, seven of these countries were in Africa, four in Asia and one in South America. In 5 out of the 10 sites for which data were available, educated men had significantly higher risk behaviour (2 showed a negative relationship and 3 were insignificant). Higher education levels of men were associated with a higher use of commercial sex workers in four out of seven sites and in most sites increases in men's education were associated with more non-regular sexual partnerships. The authors did, however, note the limitations in the data, showing the difficulties in explaining inter-country variations in prevalence because of the

importance of many cultural and ethnic differences. The data are also fairly old, and as will be discussed in this and later chapters, the relationship between socio-economic status and risk behaviour is likely to change as the level of infection and knowledge increases.

Studies of the determinants of sexual behaviour in the context of HIV provide valuable insight into the role of various socio-economic factors and contribute to the development of responses. Some researchers have attempted to go a step further and analyse the net implications of socio-economic variables on the probability of HIV infection.

1.2.4.2 Who is HIV positive?

In the literature reviewed earlier, the influence of an individual's socio-economic status on risk of infection was discussed. Cohen (1998) outlined how both poverty and wealth can lead to increased risk of HIV infection. Studies that have estimated the impact of different socio-economic status on an individual's -both men and women - risk of infection have found that risk increases, for both men and women, with an increase in socio-economic status. Ainsworth and Semali (1998) provide one such analysis and a table containing a list of other studies that have yielded similar results.

The Ainsworth study utilises data from household interviews conducted in Kagera Region, Tanzania between 1991 and 1994. Using a probit model, they compare the influence of age, education, sex, wealth and occupation on the probability of dying from AIDS, during the course of the study, with the probability of dying from another cause. The study found that more women than men were infected and that people who died of AIDS had more schooling than both those who did not die and those who died of other causes. Wealth, in the form of assets, offered protection against death from other causes, but not from AIDS. The authors noted that, because mortality data were used, the results displayed the role played by these variables some years earlier when these individuals were infected. To counter the possible response that the findings may not be relevant to current patterns of infection, Ainsworth and Semali cite the then most recent study on

socio-economic status and infection rates. The study found a positive relationship between socio-economic status and HIV infection among women attending antenatal clinics in Dar es Salaam between 1995 and 1996 (Msamanga et al, 1996).

The Ainsworth paper and the many studies listed in it have consistently found a positive relationship between socio-economic status and infection rates. This goes counter to much of the conceptual work that theorises the effects of poverty on vulnerability to infection. It would also suggest that, if there *are* protective influences associated with increases in socio-economic status, the negative influences more than outweigh them.

A number of similar studies, producing very similar results, have been conducted. Ainsworth and Semali provide a table listing these. It is, however, worth examining the findings of a more recently published study (Auvert et al, 2001), as it takes a slightly different approach.

The Auvert paper's stated objective is to 'identify factors that could explain differences in the rate of spread of HIV between different regions in sub-Saharan Africa'(pS15). The study involved the collection of data in four African cities on 1000 men and 1000 women, aged 15-49 years, in each city. The four cities were selected as representative of two high prevalence settings (Kisumu, Kenya and Ndola, Zambia) and two low prevalence settings (Cotonou, Benin and Yaounde, Cameroon) The data collection involved a socio-demographic questionnaire and tests for HIV and a number of other STIs. The researchers analysed the data with the use of logistic regressions and concluded that differences in sexual behaviour were outweighed by differences in the levels of male circumcision and STIs in explaining the variance in prevalence between the four cities. This result is of more relevance to the following section, which examines what the determinants of group prevalence are. The study did, however, also include an analysis of individual risk factors.

The individual level analysis involved univariant and multivariant logistic regressions on men and women separately in each of the four cities. The paper presents the results for

variables that were associated with HIV infection in at least one of the cities. For men, the analysis found the following to be positively associated with infection: older age, having a job, higher lifetime number of sex partners, signs, symptoms or presence of an STI, married or ever married and, in one city, not being circumcised. For women: one or more non-spousal partners in the past 12 months, higher number of lifetime sexual partners and signs, symptoms or presence of STIs.

While education is included in the analysis, the results are difficult to interpret. Given education's typically high correlation with income, analysis relating to income is particularly interesting. In total, the results of 12 education variables are reported and only one is significant: secondary education in one city is associated with lower risk. Education is expected to influence risk primarily through behaviour. If, as is the case here, behaviours are included in the regression, education would not be expected to have an influence on risk. It would be interesting to see if the education variables could help explain the behaviours reported. This was not, however, part of the examination reported on in the paper.

1.2.5 Factors influencing group prevalence

Over (1998) provides the most detailed analysis of group determinants of infection. He attempts to explain the complex issues of what leads to such large variations in HIV prevalence between countries. He argues that, although individual sexual behaviour is complex, it is reasonable to assume that economic, social, and cultural conditions shape behaviour. Based on this assumption, Over conducted regressions using urban high risk and urban low risk HIV prevalence, as reported by the US Bureau of the Census, as the dependent variables. The full data set included 50 countries, 17 from sub-Saharan Africa, 15 from Latin America and the Caribbean, 14 from Asia and 4 from the Middle East. Over used a number of societal variables as the explanatory variables. The explanatory variables included: GNP per capita, gini coefficient of the country, male and female literacy rates, percentage of the population who are Muslim, the age of the epidemic, foreign born residents, size of the military and the male:female ratio.

The selected variables yielded a good fit, all having sign that Over argued was expected and six coefficients significant at the 5 per cent level. Although the results for all the variables are of interest, of primary importance for this paper are the results for the income and inequality variables. GNP per capita of a country exhibited a negative relationship to HIV prevalence of the urban high-risk population of that country, while inequality displayed a positive relationship. Richer countries appeared to have lower rates of HIV infection among their high risk groups, while a country with a given income has higher rates of HIV the more unequal the income distribution.

Included in the Over analysis were many of the countries in which individual level studies had been conducted and had found a positive relationship between income and HIV infection. The Over study, however, appears to contradict, in terms of both his expectations and results, not only individual level studies, but other group level analysis. Stillwaggon (2000) analysed the variation in HIV prevalence between Latin American countries. She found a strong positive relationship between HIV prevalence and real GDP per capita. Stillwaggon did, however, argue that if this analysis were repeated in five years she would expect the sign to be reversed. Her reasoning for this conclusion is similar to that offered by the World Bank (1998) to explain the apparent contradiction between the Over analysis and the studies of individual risk.

1.2.6 Explaining the contradiction

The World Bank (1998) presents a collection of individual data that show a positive relationship between HIV and socio-economic status. They go on, however, to argue that most studies at the individual level were conducted in the early 1990s and that infections and deaths analysed in them occurred in the mid 1980s. At that time, knowledge of how to protect against HIV was scarce; therefore, those of higher socio-economic status, who travelled more and had resources to maintain multiply partners, were most at risk of infection. As knowledge has spread, however, they argue that the more educated and those with greater access to information would have been the first to change their

behaviour. This explanation would suggest that the individual level studies were measuring the relationship between socio-economic status and HIV infections that existed when knowledge of protective options and risks was low and that the Over study is measuring the current relationship. This explanation would also explain the apparent contradiction between Stillwaggon (2000) and Over (1998). Stillwaggon was measuring the role of income in Latin America where the epidemic was in its infancy, whereas Over included a number of countries where the epidemic had reached an advanced stage. In an effort to support the argument that the role of income changes as the epidemic matures, the World Bank presents some more recent individual level data where the positive relationship between income and HIV infection is much weaker.

The argument put forward that the role of income reverses as the epidemic matures is useful, but as yet not fully analysed. It is not clear if this is to suggest that in a mature epidemic the relationship between income and risk of HIV infection is always negative. If this were the case the highest rates of infection would be seen among the "poorest of the poor". There is currently no evidence that this has become the case. It may be the case that at some levels of income the relationship is positive, while at others it could be negative. Furthermore, has the reversal been the same, and for the same reasons for men and women? It is these and other similar questions that this thesis intends to address. To this end, theories developed in the health economic literature relating to the relationship between health, more broadly than HIV/AIDS, and income will be drawn on and applied to the questions at hand.

There exists a wealth of health economics literature examining the different effects that socio-economic status has on health more generally. The literature reviewed here relating to the role of income in the spread of HIV has not made use of recent theoretical developments in the field. A brief examination of the broader health economics literature regarding income will highlight a number of different approaches and theories that have been used to explain variations in health. These will, in later chapters, be adapted for use in analysing the problem at hand, to develop a more comprehensive theory of the role or roles played by income and socio-economic status in the spread of HIV.

1.3 Income and the health economics literature beyond AIDS

The objective of this section is not to engage in the debates in this literature. Rather it is to highlight the approaches used in explaining the role of income in generating variation in health not associated with HIV, so that they may be applied, in this thesis, to the HIV debate.

Discussions on the role of income in determining health and mortality are not new; they date back more than 200 years in the public health literature (Deaton, 2001a). The finding of a strong negative relationship between income and health, typically measured by mortality, has been attributed to a variety of causes: income having a direct influence on health, the education associated with higher income also being associated with better health. Recent debate has centred on the questions of inequality and relative income. There is wide agreement that the absolute level of income of an individual is important in determining their health. The debate, however, is about whether or not the income of others in the same group is important in determining an individual's health. Wilkinson (1996) argues that the absolute level of income is only important in determining health at low levels and that it is income inequality that is important in explaining the variation in health of wealthier countries. A number of arguments have been put forward to explain how the income of others and income inequality can affect individual health and population health. Some of the arguments stress a causal relationship between inequality and health and others a simple correlation.

The Preston Curve (1975) displays the cross-country relationship between life expectancy and income per capita (Deaton, 2001a). It shows that as average income increases life expectancy increases, but at a decreasing rate. If, at an individual level, there is a similar non-linear relationship between income and health, then a country with a given average income would have lower average health the greater the level of inequality (Deaton, 2001b). This is referred to as the individual income interpretation or the absolute income hypothesis: the curvilinear relationship between health and income is argued to be great

enough to result in the observed health variations between populations with similar average income but different income distributions (Lynch et al, 2000).

Others have argued that inequality itself is unhealthy: the pure income inequality hypothesis states that inequality reduces the health of all members of the population (Eibner and Evans, 2001). Highly unequal societies may lack 'social capital', resulting in high levels of stress and crime, which both decrease average health status (Eibner and Evans, 2001). Furthermore, inequality may lead to the under-provision of public goods that have a positive affect on health (Miller and Paxson, 2000). This is similar to the neo-material interpretation: an unequal distribution of income is a result of a number of factors that also influence the nature of public infrastructure i.e. education, health services, housing etc., all of which affect health (Lynch et al, 2000).

The most controversial theory is the relative deprivation hypothesis: inequality affects the health only of those at the lower end of the distribution, while the relatively wealthy are unharmed (Eibner and Evans, 2001). The argument states that individuals compare themselves to others within their reference group and that those who fare poorly suffer psychological stress leading to disgruntlement, possible depression, heart disease, high blood pressure etc (Wilkinson, 1997). The health of those at the lower end of the distribution can also be negatively affected by poor health-related behaviours, such as smoking and alcohol abuse. This theory differs from the absolute income hypothesis in so far as individuals may have a high absolute income, but can still be relatively deprived if their peers have even higher incomes.

The theories on how inequality affects health present a difficult analytical challenge: the definition and identification of an individual's reference group and levels of appropriate aggregation. The definition of a reference group or the level of aggregation is largely determined by the pathway through which inequality is expected to influence health. If it is assumed that inequality affects health via the under-provision of public goods then the appropriate aggregation is at the level of public good provision (Eibner and Evans, 2001).

If, however, the relative income hypothesis is being investigated there is little direction on the identification of the appropriate reference group. People may have a range of such groups, comparing themselves to colleagues, neighbours, people they read about in newspapers or see on the TV (Deaton, 2001b). Given the difficulties, researchers have attempted to define groups in a way that enables them to assume a high relevant to irrelevant ratio (individuals who are members of the reference group to those who are not). To this end such variables as state, race, education and age have been selected and then used as proxies for the unobserved true reference group (Eibner and Evans, 2001; Deaton, 1999).

The debate as to which, if any, of these theories is correct continues and indeed if they are mutually exclusive. The empirical evidence is mixed and unclear. It is not necessary, however, for this thesis to enter into this debate; it is the approaches and not the outcome that are relevant. For further discussions on the debate, the outcome and the existing evidence, see Deaton 2001a and 2001b.

The differentiation between absolute and relative income and the concept of reference groups can be applied to the HIV problem. The next chapter examines how such an application can further explain and help lead to a better understanding of the relationship between income and HIV infection. The chapter will also draw on another, very closely related, section of health economic literature. This literature examines the importance of neighbourhood in the determination of health.

There are three theories explaining why health and neighbourhood may be correlated; two suggest a causal link and one a non-causal correlation. The first is referred to as the contagion perspective (Ross, 2000). This theory suggests that people's behaviour is influenced by those around them (Crane, 1991). Health behaviours, therefore, spread and are shaped by those around an individual. The second theory, the structural perspective, suggests that residents of any neighbourhood are presented with a set of opportunities and constraints that shape their behaviour (Aneshensel and Sucuff, 1996). For example, a

wealthy area may have bike paths presenting an opportunity for safe riding (Ross, 2001). In opposition to these theories suggesting a causal link, is the compositional hypothesis, which argues that certain types of neighbourhood attract certain types of resident and the relationship between neighbourhood and health is the relationship between the individual characteristics of residents and health. The first two theories suggest that where an individual lives influences their health in terms of opportunities available or the influence of others on individual behaviour. The third theory suggests that people with similar socio-economic status, and therefore similar health, tend to live together. Again it is not relevant which theory is correct or if each has some relevance but only the idea of examining how neighbourhood characteristics may shape behaviour or how people with like characteristics live in the same neighbourhood. In many ways a similar approach has already been used with regard to HIV: discussions on dominant models of masculinity and the impact of these on risk are an obvious example (Rivers et al, 1999).

By drawing on the wealth of debate and conceptual development relating to the links between income and health it is hoped that progress can be made in explaining the links between income and HIV infection.

1.4 Summary

HIV is spread primarily through sexual intercourse. An individual's risk of becoming infected is determined by the number of times they come into contact sexually with an HIV positive person and the risk of being infected per contact. A great deal of literature and a number of empirical studies have identified numerous factors that can influence the number of contacts or the risk per contact and occasionally both. It is essential that these factors be identified to inform the development and targeting of interventions. The relationship between most of the factors identified and the spread of HIV is clear, the role played by income, and more generally socio-economic status, is, however, unclear.

Studies at the individual level have found that socio-economic status has a positive relationship to the risk of infection, mainly as a result of men with higher status and

income having greater access to sex, placing themselves and their partners at risk. Many things, however, suggest that income and risk of HIV infection should be negatively related: better access to health services and information for example. Indeed, recent studies at the country level have found such a negative relationship. In an effort to explain this apparent contradiction, researchers have argued that the individual level studies are identifying the relationship that existed in the early eighties, and that currently in advanced epidemics where information is widely available the relationship is reversed. This explanation provides part of the story; it also, however, leaves a great deal still to be explored. For example, it is unclear if the risk associated with income and access to sex has been completely negated as a result of knowledge of how to protect against infection, leaving the poorest of the poor at greatest risk, or if there are income levels where an increase in income and the associated increase in sexual partners can still result in an increase in the risk of infection. The explanations provided also do not address how the changing level of knowledge may affect the role of income differently for men and women.

In an effort to better understand the relationship between income and the risk of HIV infection this paper aims to use approaches developed in the health economics literature to examine the role of income, income inequality and neighbourhood in determining health status, not specifically HIV/AIDS.

Recent debate in the health economics literature has occurred on the possible pathways through which income inequality can influence health. One of the theories centres on the concept of relative income as distinct from absolute income, where absolute income refers to an individual's actual income level and relative income to their income relative to their reference group. These concepts of absolute and relative income and reference groups will be applied to the problem of identifying the role of income in the spread of HIV. They will be used in conjunction with theories on how neighbourhoods can influence health-related behaviours with the aim of developing a clearer understanding of the role of income in the spread of HIV.

Chapter 2

The Theory of the Dual Role of Income in the Spread of HIV

2.1 Introduction

The previous chapter outlined the literature that seeks to explain the causes of variation in the spread of HIV between populations and the distribution of infections within them. Many important studies have been done, highlighting a number of important risk factors and thereby contributing to the development of interventions. The studies which deal with socio-economic status and, more specifically, income appear, however, to contradict one another. Studies of individual risk of infection suggest that the risk increases as an individual's socio-economic status increases. Studies of differences in population prevalence levels and theory, in contrast, suggest that populations with lower incomes are at a greater risk of infection. Work has been done that goes some way towards explaining this contradiction. Although this is useful and insightful, it still leaves some questions unanswered.

This chapter applies the concepts of absolute and relative income, reference groups, and neighbourhood influence on behaviour, to the question of what role income plays in the spread of HIV. These concepts have been developed to deal with a different set of questions, specifically those relating to the variation in health between wealthy countries. They are, however, also useful in this context. By differentiating between an individual's absolute and relative income it becomes possible to separate the conflicting roles of income identified in the preceding chapter. The theory developed here suggests that higher absolute income offers protection against infection, while individuals with greater relative income (income relative to other members of their group) have a higher probability, relative to other members of their group, of being infected with HIV.

Absolute income refers to the actual income level independent of other individuals or groups. Relative income refers to the level of an individual's income compared to other members of their reference group. To make this distinction it is necessary first to define a reference group in order to facilitate the identification of variations in income within, as opposed to between, groups.

A reference group in this context is, in theory, simpler than that discussed in the previous chapter. In chapter 1, relative income was identified as a health hazard to those with low levels, mainly as a result of relatively deprived individuals making negative comparisons between themselves and other members of their reference group, resulting in stress and ill health. Deaton (2001a) points out that individuals may compare themselves to anyone, from the people they see on the TV and those they read about in the paper to their neighbours and friends. The negative influence of relative deprivation need not involve contact with other members of an individual's reference group. In such a context, identifying an individual's reference group and their position in it is obviously very difficult if not impossible.

The theory developed in this chapter, however, examines the role of income in the sexual transmission of HIV, which obviously requires physical contact. Although some people have television stars as sexual partners, this is not widespread and, provided there is no physical interaction, the possibility that an individual might compare themselves to TV stars is irrelevant. In this context it becomes far simpler, although still not easy, to define an individual's reference group. An individual's reference group in relation to a study of HIV infection can be defined as the group of people from which an individual selects sexual partners, plus those from whom their potential sexual partners select their sexual partners. That is to say that an individual's reference group comprises both potential sexual partners and those who are selecting partners from the same group. In the heterosexual context, an individual's reference group includes both men and women. Although this is a simpler reference group to identify than the comparison group discussed earlier, there are still some practical difficulties, which will be addressed in detail in the next two chapters. With the definition of the reference group it is possible to

define relative and absolute income more clearly. An individual's absolute income is their income independent of their, or any other, group and their relative income is their income relative to other members of their reference group.

Building on the above definitions, this chapter develops a theory on the role of absolute income and relative income. It deals firstly with each independently; then they are examined in combination. Based on the understanding of the underlying roles of relative and absolute income, an attempt is made to hypothesise what the observed relationship between HIV prevalence and income should look like.

The relationships described relate to the sexually active population. Within the sexually active population there are substantial variations in HIV prevalence between age groups; the pattern of infection rates observed is fairly standard throughout the sub-Saharan region (Stover, 1999). While some of the age variation may be a result of income, it is not discussed in detail at this stage, as it would complicate the arguments a great deal. In later analyses, the effects of age are controlled for to avoid misinterpretation of results. This analysis concentrates on life time risk, that is the risk of ever being infected, rather then how income affects the timing of infections, which may well be an important issue, but is beyond the scope of this thesis.

The following sections discuss the role of HIV in relation to a mature epidemic. Following this is an examination of how the relationship might vary at different stages of the epidemic. The analysis is based on a static approach. Reference to an increase in income in this context refers to the risks associated with individuals or groups with higher income compared to those with lower income, rather than the change in risk associated with an individual's income increasing. The dynamic analysis of changing income levels is presented in chapter 5.

2.2 The role of absolute income in the spread of HIV

2.2.1 Individuals' absolute income

An individual's absolute income, as defined above, is their actual level of income independent of others in society or their reference group. As absolute income increases, a number of factors influence an individual's risks of becoming infected with HIV. The following section outlines the path from changing absolute income to changing risk of HIV infection.

Firstly, as absolute income increases an individual's access to health services increases; this reduces risk, primarily through the treatment of STIs. Better access to health services leads to better treatment and lower rates of STIs. Lower rates of STIs result in a lower probability of infection per sexual contact with an HIV positive partner, thus reducing the overall probability of infection. Returning to the simple model in chapter 1, this can be shown mathematically by combining the literature outlined in chapter 1 with the concept of absolute income as a factor independent of relative income:

$$\text{Prob(Individual HIV)} = \text{Contacts} * \text{Probability} \quad (1)$$

$$\text{Probability} = f(\text{sex-behaviour}, \text{STI1}*\text{type}, \text{STI2}*\text{type}, \text{condom}) \quad (2)$$
(where Probability $\propto$ STI1&2, 1/condom)

$$\text{STI1} = f(\text{Access}, \text{HSB}, \text{condom}) \quad (3)$$
(where STI1$\propto$ 1/Access)

$$\text{Access} = f(\text{AY}, \text{AHS}) \quad (4)$$
(where Access $\propto$ AY)

$\therefore$ probability $\propto$ 1/AY

Where:

Prob(Individual HIV) is the probability that an individual will become HIV positive

Contacts is the number of times an individual comes into contact with HIV

Probability is the probability of infection per contact

Sex-behaviour is the type of sexual intercourse (in the heterosexual context this refers to vaginal, dry, anal or oral sex)

STI1 is the presence of an STI or not for the individual

STI2 is the presence of an STI or not for the partner

Type is the type of STI present (different STIs increase the probability of infection to different degrees)

Condom is the frequency of correct condom use

Access is access to health services

HSB is health seeking behaviour

AY is absolute income

AHS is availability of health services.

The model shows the link between increases in absolute income, through access to health services, and the resultant reduction in probability of transmission, leading finally to a reduction in the risk of the individual becoming infected with HIV. The factors that directly increase the probability of being infected are the sexual behaviours, prevalence of STIs and condom use. But the prevalence of STIs is a function of health seeking behaviour and access to health services, which in turn are related positively to income. Therefore as income increases, access to health services improves, STI rates fall; as a result the individuals risk of infection decreases as absolute income increases. The model, however, also brings in a number of other factors, particularly the importance of health seeking behaviour and of the partner's STI status. These and other factors are discussed later in the chapter and are shown to be very important in determining the ways in which absolute income influences risk of infection.

An increase in absolute income also increases individuals' access to information. Individuals with greater income will have more access to all types of media. This access to information increases their awareness of HIV and the ways in which to protect themselves against infection. The use of condoms, for example, would be expected to

increase with income. As income increases access to information on how condoms can protect against the chances of infection, so does knowledge of the consequences of infection. Given a constant cost of purchasing and using a condom, the increase in perceived benefits of use should increase the probability of use, therefore decreasing the probability of infection.

Information on protective means can be better understood with increased education. Education and income are typically correlated; therefore higher levels of absolute income would be correlated to greater levels of understanding and therefore greater perceived benefits of adopting safer sex practices and greater perceived costs of not doing so. Filmer (1998), in the study outlined in the previous chapter, found that the probability of using a condom with a non-regular partner and education are strongly correlated.

In the context of the simple model outlined above, better access to information would be expected to improve health seeking behaviour (equation 3), lead to safer sex behaviours and higher condom use (equation 3) and therefore lower rates of STIs (equation 3). All of these influences would reduce the probability of an individual becoming infected, by reducing both the number of contacts and the risk per contact (equation 1).

In addition to the factors outlined above, three other factors related to the level of absolute income, may influence the probability of an individual becoming infected, but the relationships are not as clear. Increases in absolute income improve economic and food security and increase individuals' ability to travel.

As absolute income increases, individuals' economic security improves. In the literature on factors influencing the spread of HIV, coping mechanisms associated with economic insecurity are associated with greater risk of infection. The most obvious of these are commercial sex work and migrant labour. Commercial sex work is, however, a function of both an individual's absolute income and their income relative to others and will be discussed later. Migrant labour is a function of an individual's absolute income and the

geographical distribution of income is an issue that will be discussed in more detail with regard to relative income.

These complexities associated with relative income aside, it would be expected that an increase in absolute income would lead to greater economic security and therefore reduced risk behaviours linked to economic survival. Two factors, however, cloud this apparently simple relationship. Firstly, if the increase in economic security is a result of a coping strategy, it is unlikely to lead to safer behaviour: for example, an increase in economic security resulting from a woman entering the commercial sex work industry. Secondly, the risk linked to these strategies may not always exhibit a linear relationship to income. For example, a subsistence farmer may not have sufficient security to obtain financing to increase output and sell her surplus. An increase in economic security may allow her to obtain finance, increase production and travel to the market to sell. On the one hand, this increased travel may expose her to greater risk, as the travel exposes her to a larger sexual network; on the other hand, the additional cash may mean that she no longer has to exchange sexual favours for essential household goods. Although the conclusion is not clear-cut, on balance, it seems fair to assume that an increase in economic security, brought about by an increase in absolute income, reduces individuals' risk of infection by reducing the probability of adopting coping strategies that would increase the risk.

The second, but as yet unproved, factor is the effect of improved diet, resulting from a greater absolute income, on an individual's risk of infection. Stillwaggon (2000) found a negative relationship between calorie intake and HIV prevalence. If this relationship is causal, then an increase in absolute income that increases food security, and as a result leads to a stronger immune system, would lower the probability of HIV infection per sexual contact with an HIV positive partner. Stillwaggon discusses a number of different dietary-related problems that could affect transmission. A lack of vitamin A, for example, weakens the vaginal wall increasing the probability of tearing during intercourse, resulting in bleeding and increased probability of transmission. These links have never been medically proven with regard to HIV, but there is not even a suggestion

that improved diet and food security could *increase* an individual's risk of infection. If the relationship exists, it strengthens the negative relationship between absolute income and risk of infection associated with the other factors mentioned; if it does not exist, it does nothing to weaken the relationship.

Thus far, all the factors associated with absolute income that have an influence on individuals' risk of infection appear to reduce the risk. The only variable that changes as a result of changes in absolute income that could be contrary to this trend is the increased mobility associated with a rise in absolute income.

As individuals' absolute level of income increases, they have more access to a wider variety of modes of transportation, ranging from the ability to buy a bicycle to purchasing plane tickets, or even an aeroplane. Increased means and more ability to travel would suggest a wider range of potential sexual partners, possibly resulting in increased risk of infection for individuals or groups with higher absolute income. The relationship is, however, clouded by a number of complicating factors: the reason for travel, mode of transport, length of stay, family accompaniment and sexual networks.

An increase in absolute income may improve individuals' access to transport, but it may also change their motivation for travel. Individuals at lower levels of absolute income may travel to find work, passing through many areas in their search, while an individual with higher absolute income may travel for holidays or brief business trips. These differing motivations expose individuals to different types of risk and it is not immediately clear which type of risk is the greater.

The issue is further complicated by the possible differences in modes of transport. An individual with high absolute income may chose to fly, whereas an individual with low absolute income may be forced to travel by taxi. The taxi journey will take longer, exposing the individual to more people.

Furthermore, lower absolute income may make return journeys infrequent; a migrant worker with a low level of absolute income may only return home a few times a year, which increases the probability of sexual involvement while away from home. High levels of absolute income provide individuals with the means to travel with their family, possibly reducing sexual contacts with new partners.

Individuals with high absolute income, who travel alone, are more likely, compared to individuals with low absolute income, in their travels, to select sexual partners with high absolute income than are travellers with low absolute income. Their partners are, therefore, less likely to be infected, thereby lowering the individual's risk. These complicating factors make it difficult to conclude what effect an increase in ability to travel has on an individual's risk of infection.

With the exception of the relationship between absolute income and travel, all other factors linked to absolute income either reduce the risk of an individual becoming infected, by reducing the number of contacts or the risk per contact, or have no effect. It would appear reasonable, therefore, to hypothesise that as the absolute income of an individual increases so the risk of being infected with HIV decreases.

2.2.2 Average group absolute income

Examining the role of absolute income with reference only to the individual is an oversimplification. Many of the factors outlined above are shaped not only by the individual but also by their reference group. The role of absolute income needs, therefore, to be discussed in terms of group dynamics. This is a distinct discussion, from the role of relative income, in so far as it will examine the importance of average group income in absolute terms as opposed to individual incomes within a group compared to the average of that group.

This section argues that that there is a strong group effect in the realisation of protection derived from absolute income. Average absolute income of the individual's reference

group may be as, if not more, important in determining the level of protection derived, as the individual's level of absolute income. Much of this argument is based on the assumption that there is a high degree of overlap between an individual's reference group and the community in which they live. Although the existence of migrant labour and a number of other factors mean that sexual networks, and therefore reference groups, are not so clearly geographically defined, it seems reasonable to assume that the majority of sexual partners are drawn from an individual's community.

Access to health care, identified as a protective pathway through which an increase in absolute income reduces the individual's probability of infection with HIV, is influenced not only by an individual's absolute income but also by the availability of health services and their health seeking behaviour. The availability of health services to an individual, is to a large extent, determined by those available within their community. These services are a function of, among other things, the average absolute income of the community (in this context a geographic area) in which they operate. Other things being equal, a community with a higher absolute income will have better health facilities; therefore, an individual living within such a community would have greater access to health services.

Furthermore, the health seeking behaviour of an individual is arguably a function of their own education, access to information, and the normal behaviours of their reference group. As was noted in the literature in the previous chapter, behaviours, according to both the contagion and the structural perspectives, are arguably shaped by the neighbourhood in which individuals live. If this is the case regarding health seeking behaviours, then again the average income and education of the group, assuming that education and income are important in determining health seeking behaviour, play a role in determining the probability of an individual becoming infected, by influencing their health seeking behaviour.

The above behavioural context argument can also be applied to the protective influences associated with information and sexual behaviour. If the context argument is combined with the argument outlined earlier for an individual that absolute income, education,

access to information and level of protective behaviour are linked, then the probability of an individual being infected is further influenced by the average absolute income of their reference group, which is made up primarily of members of their community. This is the case because as the average level of access to information increases the more likely it is that 'normal' behaviours, based on information accessed, and acceptable to the reference group, will be protective. This is by no means a new argument. Rivers and Aggleton (1999), among others, argue that, particularly in regard to sexual roles, individuals' behaviours are shaped by the dominant view of the society in which they live. Those who do not conform are often excluded and viewed as unsuccessful, the treatment of homosexual men in many societies being an example of such exclusion.

The above context arguments suggest that, while the level of an individual's absolute income is important, the average absolute income of their reference group is as, if not more, important. This argument is reinforced by an examination of how the health of the group influences the probability of infection for the individual. Firstly, in the simple model outlined (equations 1 – 4, at the beginning of this chapter) the presence of a STI in the partner was highlighted as a risk factor for the individual. According to the definition of a reference group, sexual partners are drawn from the individual's reference group. The higher the average absolute income of the reference group the lower the probability that an individual's sexual partner will have an STI. This factor influences the risk per contact. The protective role of increased group absolute income, however, goes further to reduce not only the risk per contact but probability of contacts with an HIV positive partner, as partners are drawn from a group with lower risk. A higher average absolute group income, therefore, reduces the risk of an individual within that group being infected, regardless of their position within that group.

The protective influence of absolute income can be approximated as the protection offered by reference group average income, because of the importance of context in the realisation of these protective influences. References to absolute income, unless otherwise stated, for the rest of the paper, therefore, refer to average reference group absolute income.

Having identified the protective factors of absolute income and the direction of the relationship to HIV infection, two other issues require consideration: the strength of this relationship and the importance of the stage of the epidemic in determining both its slope and gradient.

2.2.3 The strength of the relationship

The factors outlined in the previous sections suggest that individuals and groups with high absolute income have lower risks of infection than groups and individuals with lower absolute income. The relationship is, however, unlikely to be linear in nature. This section discusses the likely strength of the relationship at different levels of absolute income. It proposes that, at very low levels of absolute income, increases are likely to offer little protection and that this stage is followed by a phase of rapid decreases in risk, which quickly slow. This relationship differs only slightly from that commonly associated with health status and absolute income variations. The Preston Curve (1975), which was mentioned in the previous chapter, suggested that at low levels of income increases in income resulted in larger gains in health status than the same increase at a higher starting point. The difference appears only at very low levels of income where on the Preston curve the greatest benefits are achieved from increases in income, but the argument relating to HIV suggests that greater benefits are achieved at still low, but slightly higher, incomes. The foundation for this marginally different argument is based on the issues relating to access to health service and the benefits derived from such access and the decreasing importance of information.

Access to health services, although it improves with increases in absolute income, does so at a slow and discontinuous rate. The jump from primary to secondary and tertiary care is large and expensive and requires high levels of absolute income to sustain. Similarly, changes in health seeking behaviour based on social norms are not sudden movements and small changes in absolute income are likely to lead to only small changes

in behaviour. These two factors suggest that as absolute income increases, increases in access to health services are small and discontinuous.

The protective aspects of health care in regard to HIV infection are not a result of advanced medical procedures: access to open heart surgery, for example, will do little to slow the spread of HIV. The protection is derived primarily through the provision of appropriate basic health care, especially the treatment of STIs. Therefore, although increases in absolute income may lead to only small changes in access and health seeking behaviour, these small changes at low levels of absolute income may lead to large changes in risk. This would be the case if they meant the introduction or improvement of STI care. If, however, the improvements do not lead to the introduction of appropriate care, there is likely to be little benefit in terms of reduced risk of HIV infection. The benefits of improvements in access to health care are, in any event, short-lived. Based on these factors, the benefits of increased health care can be broken down into three phases: the poverty trap, basic health care, advanced care.

The poverty trap phase relates to levels of absolute income where available health care is of such poor quality that increases in absolute income, although they may improve access, do little to reduce risk of infection. For example, a community with extremely low levels of absolute income and no access to health services will have a similar level of risk as a slightly wealthier group with access to only very basic health care. This is because if the available health services do not provide services that reduce risk – for example, treatment for STIs - then there will obviously be no reduction in risk. This factor accounts for the difference between the relationship outlined here and that typically associated with health status and absolute income. While a small increase from a very low level may have large returns in terms of improved average health status (the basis of the common relationship), if they are not relevant to HIV they will have little impact on risk.

The basic health care phase refers to the stage where increases in absolute income lead to the introduction of the levels and types of care that offer protection against infection,

most notably treatment for STIs. During this stage small improvements will lead to large changes in risk.

The final phase, advanced medical care, refers to the stage where basic health care has approached its maximum and increases in absolute income improve access to advanced medical treatments. Although these treatments may have many other benefits, in the context of the spread of HIV, improvements in this phase will do little, to reduce the spread of HIV. Advanced medical treatments can, however, prolong the life of those infected, this however is a distinct issue.

The relationship between the phases and the average risk of infection for members of groups can be depicted graphically as follows:

Figure 2.1 HIV risk and access to health services

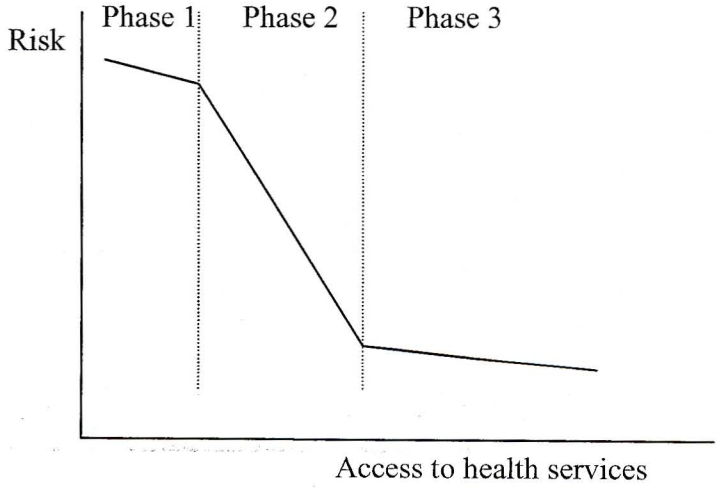

The access to information is likely to display a similar pattern, to that displayed above. It is not simply the knowledge of how to protect against HIV that protects against HIV, but the behaviour based on that knowledge. Again, changes in absolute income are likely to lead to only small changes in behaviour but, as with health care, a small change in behaviour can lead to a substantial change in risk. The similarity to health care continues in so far as, after the provision of basic information and the behaviour change based on that, there is little protective benefit from additional knowledge. The basic information provides the knowledge that sexual intercourse is the primary route through which HIV is transmitted and that condoms reduce the risk of transmission reduces risk a great deal. Additional knowledge of virus subtypes and viral load, for example, may lead to further slight changes in behaviour, but the benefit is not comparable to the fundamental behaviour change occasioned by the basic information. The returns on information, as with health care, exhibit diminishing marginal returns.

Based on the above assertions, the relationship between individual and group risk of HIV infection and absolute income can be displayed graphically as follows; as more effects are considered the relationship depicted is smoother than in the health example which considered only one relationship:

Figure 2.2 The relationship between absolute income and the risk of HIV infection

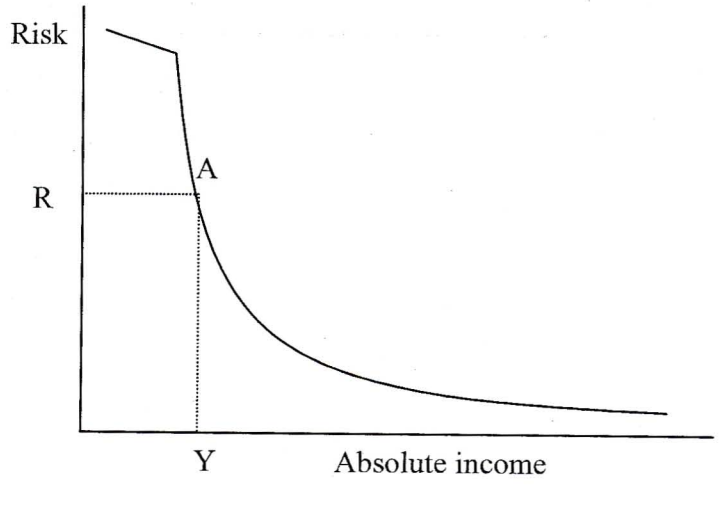

Point A displays the relationship between the risk, R, of HIV infection for an individual member of a reference group with absolute income Y. At low levels of absolute income the risk of HIV infection reduces slowly, as behaviour change is slow and appropriate health services are not available. This is followed by a rapid reduction in risk as small changes in behaviour and health services lead to major reductions in risk. As diminishing marginal returns to income set in, the reduction in risk for a given increase in absolute income becomes smaller as the improvement in health services becomes less and less relevant and the increased information less and less useful for behaviour change. The representation assumes that the rapid change from health services and information occurs at a similar level of income. Those with access to basic primary health care are likely to have access to basic health information, if only from the providers of such basic health care, whereas groups of people with no access to health services are unlikely to be well-informed on protective practices.

The relationship displayed has been described and developed in reference to a mature epidemic. Given the importance attached in the literature to the time and stage of the epidemic in relation to the role of income, it is essential to discuss how this relationship might differ at the various stages of the epidemic.

2.2.4 The onset of the epidemic

Before an HIV epidemic can begin in a population the virus has to be introduced. Prior to this introduction, nobody faces the risk of being infected regardless of sexual practices of STD infection. The simple model outlined at the beginning of this chapter shows how zero contacts with HIV infected partners lead to zero risk. At this point, the only people who are at risk are those who travel and mix sexually with populations, other than their own, where HIV is present.

As noted previously, an increase in absolute income increases an individual's ability to travel. This travel puts them into contact with other population groups where they may come into contact with HIV. At the very onset of the epidemic there is clearly a positive

relationship between absolute income and risk of infection. Figure 2.3 represents the situation:

Figure 2.3 Absolute income and risk at the onset of an HIV epidemic

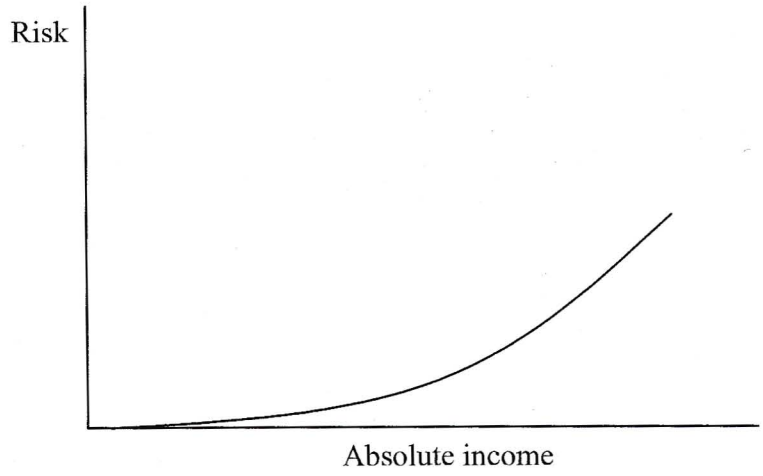

The relationship depicted in figure 2.3 is reinforced by the lack of protection offered through absolute income as a result of knowledge. At the onset of the epidemic, individual and group knowledge is low and sexual behaviour unchanged. The one protective influence still favouring those with higher absolute income is their better access to health care. Better access to health care and lower rates of STIs mean that individuals from groups with high absolute income have a low risk of infection per contact. At the very onset of the epidemic, when individuals with low income have sex only with other individuals in their own population where there is no HIV, those with high absolute income are still at greater risk. This greater risk will be maintained until the probability of low income individuals coming into contact with the virus reaches a stage where their increased risk per contact places them at greater overall risk of infection. HIV is introduced into the population by those with high levels of absolute income and increasingly by those with lower levels, as the distance to populations already affected by HIV becomes shorter. As information becomes more widely available and the epidemic matures, the relationship between absolute income shifts from that depicted above to that displayed in figure 2.2. This is similar to the argument proposed by the

World Bank, 1998) as to why individual and country studies of risk and prevalence differed in their findings relating to income.

2.2.5 Summary of the role of absolute income

In theory, higher levels of absolute income offer protection against infection with HIV. This protection is offered via a number of pathways. Firstly, higher absolute income means improved access to health services. Such improved access reduces the prevalence of STIs, thereby reducing the risk of HIV infection per sexual contact with an infected partner.

Secondly, greater absolute income leads to better access to information. This, coupled with higher rates of education, results in better-informed sexual behaviour and lower risk of infection.

Arguably, higher absolute income reduces the risk of infection by improving economic security and possibly by improvement in diet.

Increased absolute income does, however, increase ability to travel. At the initial stage of the epidemic this greatly increases risk of infection for those with high income. At later stages, the influence of this factor is not clear and is very unlikely to be so strongly positive as to reverse the net negative influence on risk resulting from changes in absolute income.

Group influence on behaviour and the importance of available resources in determining behaviour suggest that the protective aspects of absolute income work at a group level. This suggests that it is the average absolute income of a group that is important in determining the group's risk and the risk of individuals within it.

Conceptually, increases in absolute income offer little protection at very low levels of absolute income. This situation is, however, short-lived and increases of absolute income

at low levels should quickly change and begin to have a substantial influence on risk of infection. Diminishing marginal returns is, however, likely to set in quickly and returns to increases in absolute income, in the form of reduced risk, are likely to be very small at high levels of absolute income.

The above discussion has dealt in detail with the role played by absolute income in the spread of HIV. A number of issues that were identified in the literature are not mentioned here. Most notably, there is no reference to the greater access to sex associated with higher income and no mention is made of male/female income differentials. These exclusions were made because they are factors associated with individuals' income relative to others and not their absolute level of income. It is these factors that will be dealt with in the following section and it is the ability to separate them from the factors discussed in this section that makes the use of the absolute and relative income concepts so useful.

2.3 Relative income

Relative income considers an individual's income relative to other members of their reference group, in other word an individual's income relative to the average income of the group. An individual may be a member of a group with a very low average level of absolute income but, because they earn a little more than that average, they are considered to be relatively wealthy. Similarly, an individual who is a member of a reference group with an extremely high average level of absolute income may be classed as relatively deprived if they have lower than the group average income.

The pathways through which relative income influences an individual's probability of being infected with HIV are very different from those associated with absolute income. The protection offered by increases in absolute income, according to the theory developed in this chapter, is a result primarily of two factors - access to health care and information. These two protective factors are shaped by the characteristics of the individual's reference group. Access to health care is, as described earlier, a function not

only of an individual's own absolute income, but of the facilities available to their group and the dominant behaviours relating to health-seeking behaviour. Similarly, increases in access to information provide protection against HIV by the influence they have on behaviour. Again, this behaviour is shaped by the context in which it occurs and what is acceptable and valued within an individual's reference group.

The importance of group characteristics in the realisation of the protective influence of absolute income suggests that variations in an individual's absolute income within the group would have little influence on individual risk. This is because it is the average income of the group that determines access to health services and information and shapes social constructs of acceptable behaviour. Variations in relative income are by definition variations of income within the reference group, which means that they can be discussed separately from absolute income. In other words, changes in risk associated with different levels of individual absolute income between members of the same group are assumed to be so small that they are overwhelmed by the impacts of relative income. In other words, the average absolute income of a group sets the mean risk and relative income distribution about that mean.

It is, therefore, assumed in this analysis that individuals within groups have very similar levels of protection resulting from the average absolute income of their group. Higher relative income within the group does not offer protection against the spread of HIV through increased access to health services or behaviour change resulting from greater access to information, as this protection is realised largely at a group level. This assumption is, however, an oversimplification, as individuals may be able to select from the health services available, according to their ability to pay, thereby accessing differing levels of health care. Moreover, information may lead to some individuals protecting themselves despite the negative reaction this may illicit from their reference group. Behaviour change in the face of negative reaction is the exception rather than the rule. In any event, the benefits of slight variations in health care are likely to be small unless it means the difference between accessing STI care or having no treatment. Even in such a situation, however, it must be remembered that the individual's risk is also affected by

the health of their partner drawn from their reference group. This simplification, therefore, seams reasonable, and it can be argued that these variations in risk within groups, resulting from differing access to health services and information, are negligible and can be ignored.

Assuming that the protective influences of increased absolute income are realised largely at a group level - that is they determine the average risk of a group and are therefore negligible in relation within group changes in relative income - the influences on risk associated with relative income changes, variation about the mean, can be more closely examined. With similar access to health services and similar sexual behaviours – such as condom use – the probability that a member of a reference group will be infected, compared to another member of the same group, will be largely determined by the number of contacts that an individual has with HIV positive sexual partners. This in turn will be affected by two factors: the number of partners they have and the probability that each partner is infected. Returning to the simple model and applying the concept of relative income to the literature reviewed in the previous chapter:

Prob(Individual HIV) = contacts * probability (5)

Contacts = f(Prob(Partner HIV), No. of partners) (6)

(where contacts $\propto$ Prob(Partner HIV), No. of partners)

No. of partners = f(preferences, power, attractiveness) (7)

Power = f(rY, norms) (8)

(where Power $\propto$ rY)

Attractiveness = F(X, rY) (9)

(where attractiveness $\propto$ rY)

Where:

Prob(Individual HIV) is the probability that an individual will become HIV positive

Contacts is the number of times an individual comes into contact with HIV

Probability is the probability of infection per contact

Prob(Partner HIV) is the probability that an individual's sexual partner is HIV positive

No. of partners is the number of sexual partners an individual has in a given time period

Preference is an individual's preference for a number of sexual partners

Power is an individual's power to enforce their preferences on others individuals in their reference group

Attractiveness is an individual's attractiveness to other members of the individual's reference group

X is some set of variables considered to make an individual more or less attractive

Norms are a set of common behaviours and practices

rY is an individual's relative income.

The model depicts a situation where the number of sexual partners an individual has is determined by their sexual preferences, their power over others to realise these preferences, and their attractiveness to others. Both power relations and attractiveness are, in part, explained by the individual's relative income, as distinct from their absolute income. Within a reference group with a given average absolute income, those who are relatively wealthy are, other things being equal, more powerful and attractive because of their control and access to resources. This is not a result of their absolute income, because if they were in another reference group, with a higher average absolute income but their income remained constant, they would have less power and be less attractive because their position relative to others would be weaker. The model does not stipulate the direction of many of the relationships. On closer examination it will become apparent that these relationships are influenced by the individual's sex. The following two sections examine the direction of these relationships for male and female members of reference groups.

2.3.1 Male relative income

Sexual preferences and sexual behaviour, as stated many times, are shaped by the society or, in this case, reference group, in which these behaviours occur. The role of relative income will differ based on these preferences. In this thesis, it is assumed that male promiscuity is seen as a sign of success, power and manhood, whereas it is frowned upon for women. This is the most commonly found social construct of sexual roles for men and women (Rivers and Aggleton, 1999) and for this reason the examination of relative income will be framed in this context. Indeed, Filmer (1998) in his cross country study, details of which appear in the previous chapter, found that in all the samples he examined it was more likely that a man would report having a non-regular partner and much more likely that those who had more than one non-regular partner would be men.

Other things being equal, individuals with greater relative income have greater power over others in their reference group. The level of this power and over whom it can be exercised will be dealt with in more detail later. It is also reasonable to assume that, holding other things constant, men with higher relative incomes will be more attractive to potential partners. The increases in power come from men's income relative to women in the group, while their attractiveness comes from the income relative to other men. If men are assumed to prefer more partners to fewer, then these increases in power and attractiveness would increase the number of sexual partners an individual has, thereby increasing their risk of infection as a result of greater relative income. If sexual partners have the characteristics of a normal good that exhibits diminishing marginal utility, then as relative income increases the demand for additional sexual partners will increase at a decreasing rate. The relationship can be represented graphically as follows:

Figure 2.4 Male risk and relative income

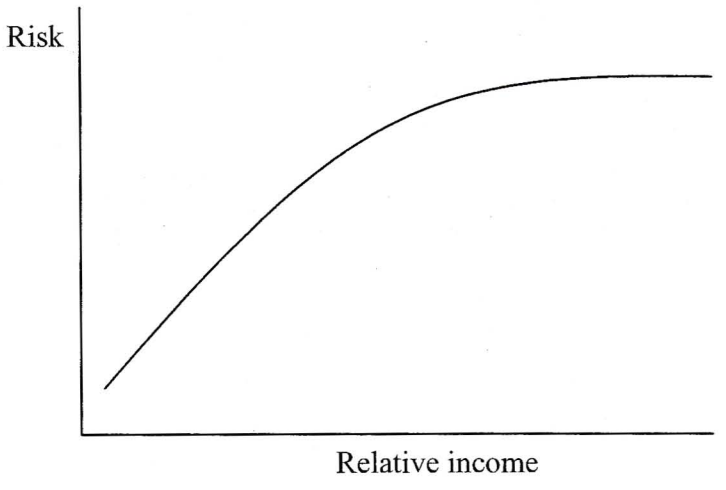

The relationship described above is not meant to reduce sexual partnerships to issues of financial resources and commodities or to suggest that all men seek to have as many partners as possible. Indeed, many men will remain with only one partner, regardless of their relative income. Similarly, many women will decline sexual relations with one man in favour of another, even if the first is relatively more wealthy. All that is being suggested is that, subject to great individual variations, as male income increases relative to others within their reference group they will, on average, have more sexual partners and, as a result, greater risk of becoming infected with HIV.

2.3.2 Female relative income

To understand the risk associated with different levels of female relative income it is necessary to further examine the reasons for the described male risk pattern. Higher levels of relative income are associated with a higher probability of HIV infection for men, as a result of their greater power within the group and their attractiveness relative to

other men within their group. Both of these factors influence the risk of women very differently.

If men with higher relative income have greater power over women within their reference group, which increases the availability of sexual partners as men have more power over the terms and frequency of sexual relations, then the greater the disparity in income between possible partners, the greater the power. Relatively wealthy men, therefore, have greatest power over the most deprived women within their reference group. This power relationship increases the risk of men at high levels of relative income and women at low levels. For example, in a group with low absolute income, relatively deprived women may be forced to rely on sexual favours to gain access to household essentials controlled by relatively wealthy men. This situation may occur both within and between households. The risk of relatively deprived women is, therefore, influenced by the risk of relatively high-income males.

The examination of the influence of power relations suggests that women with higher income relative to other members within their reference group would have lower risk of infection, compared to women with lower relative income, because the power men have over relatively more wealthy women is weaker. The second factor, attractiveness, which influences male risk, however, also influences female risk. The power relations argument is not dependent on the nature of the relationship and can be made with reference to long term and temporary sexual relationships. The argument made in relation to attractiveness relates to long term relationships and the risk associated with these long term sexual relationships may result in a very different risk pattern to that associated with power relationships.

According to the principle of positive assortative matching, a woman will partner with a man of equal or higher socio-economic status (Boulier and Rosenzweig, 1984). In the context of HIV risk, woman with higher income relative to other women in the same reference group (and therefore higher socio-economic status) will be at a greater risk of

having an HIV positive partner as a result of his higher relative income and associated higher risk.

Returning to the simple model, within a reference group, women with higher relative income face a higher probability that their long term sexual partner is HIV positive and therefore are at greater risk. The women's higher relative income, however, reduces the number of men who have power over her and therefore reduces the average number of sexual partners resulting from men having such power over her, thereby reducing her risk. The net effect of an increase in relative income is, therefore, dependent on which of these effects is dominant. While having multiple partners has been shown to increase risk of women (Auvert et al, 2001), the majority of women infected with HIV in Africa, 60-80 percent, had only one sexual partner (Adler et al, 1996). The dominant risk effect, therefore, is assumed to be the one that mirrors their partner's risk. It must, however, be kept in mind that situations may arise where the risk associated with unequal power relations may be so great that it outweighs the risk pattern associated with the partners. Conceivably, this could be the case for young women, very few of whom have long term partners, but many may be at risk as a result of the unequal distribution of wealth in their reference group. In such a case, higher relative income may, at least for the time being, reduce the risk of young women. This theory, as stated earlier, does not deal with the age differentials, but rather with the average for all individuals. In this case it is assumed that the partner's risk is the dominant effect, although this may not always be the case.

The risk associated with different relative income levels is, however, likely to differ for women for two reasons. Firstly, if women partner with men of equal or greater socio-economic status, the average risk for women at any given income will mirror the risk of men with slightly higher income; secondly, while the tracking of partners' risk is arguably the dominant factor, the risk associated with power relations cannot be ignored. Arguably, this factor increases the risk of those with low incomes and has less and less effect as income rises. Based on these two factors, the risk of women is hypothesised to track the pattern of male prevalence, but at a higher level, particularly at low levels of income. figure 2.5 depicts this relationship.

Figure 2.5 Male and female risk and relative income

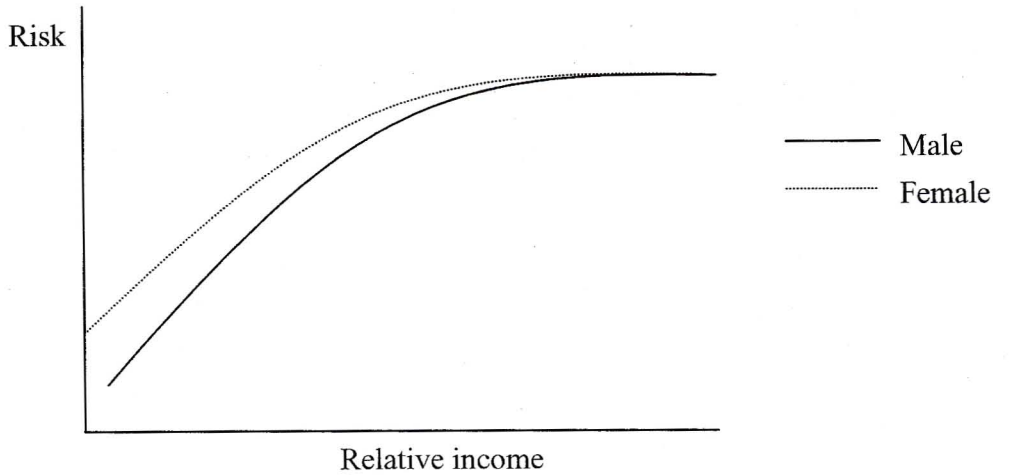

The figure above shows how, in the context of the simple model, female risk increases with relative income. At low levels, female risk is greater than the risk of males with the same level of income, because women at these low levels face risk associated with unequal power relations and the risk of their partner. As relative income increases, the difference between female and male risk decreases. Female average risk is, therefore likely to be higher because at different levels of income it is either higher or the same.

The relationship between relative income and female risk of infection, as described thus far, has not addressed the possibility of higher demand for sexual partners by women with higher levels of relative income. Unless this were more widespread than for males, so that an increase in relative income led to an increase in the number of sexual partners greater than the increase for males, this would only change the size of the gap between the two lines. If it were the case that the increase in sexual partners was greater for

women, the lines would be reversed and men would be at greater risk. This is, however, very unlikely and assumed not to be the case.

The above is presented as a theory of the average, that is the influence of relative income on the average risk of people at that level of income. There will be substantial individual variations. Some of this variation may be a result of personal moral and religious beliefs and other individual variations that do not fall neatly into the theory. There is not space or cause to discuss the impact of these, with the exception of one. Commercial sex workers present a special case. Sex workers tend to be young and relatively poorly educated (Ahlburg and Jensen, 1998). Their risk is therefore a result of their low status, but once they begin work they may have relatively high incomes. The more clients they have, other things being equal, the greater the risk and the income[2]. Although commercial sex workers would appear to have the same relationship, in terms of direction rather than magnitude, between relative income and risk, the causality is reversed. The higher risk has led to greater income[3]. The emphasis of this chapter, and the rest of the thesis, is on the risk of average men and women and on special cases only in so far as they influence the risk of the average.

2.4 Combination

Thus far, absolute and relative income have been discussed independently of each other. In the theory developed here, absolute income sets the level of group risk, the average risk of the group, while relative income explains the variation in risk within groups, about the average. The following section examines how the two work together to determine individual risk of HIV infection. The following figure represents a simple combination of the absolute and relative income affects graphically:

[2] Within in this special case there may be even more special cases. For example a high earning commercial sex worker may have more power over the terms of sexual exchange and may be able to negotiate safer sex.

[3] This is one special case there are likely to be others.

Figure 2.6 Simple combination of absolute and relative income

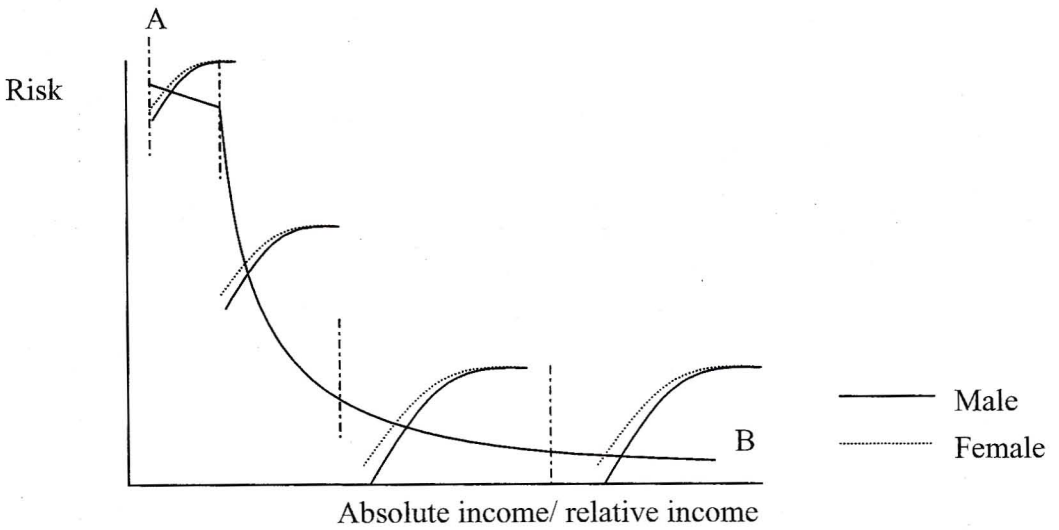

The above figure overlays the described relationship between absolute and relative income and risk of HIV infection. Line AB represents the protective effect of increases in absolute income over the entire spectrum of income levels. The vertical dotted lines represent the division of the population into reference groups and the curved lines within the reference group boundaries represent the increase in risk associated with increases in relative income - one line (solid) for males and one (dotted) for females[4]. An individual's risk is, therefore, a function of both their absolute and relative income. The figure is, however, an oversimplification and suggests that the relationship between relative income and risk, although placed in a range by absolute income levels, works independently of absolute income. That is to say that the gradient of the relative income curve is unaffected by the absolute income of the reference group. For this to be the case, the power held by relatively wealthy men over relatively deprived women would have to be the same regardless of the absolute level of income. This is, however, unlikely.

[4] In reality it would not be possible to divide the population up in to such neat and non-overlapping reference group. This representation is grossly simplified, but the simplifying assumptions will be relaxed

At low levels of absolute income, women may be forced to rely on financial and material support from relatively wealthy males for the basic means of survival. The power held by these men over these women is therefore great. In a reference group with high absolute income this power would be largely diluted; thus the gradient of the relative income curve is reduced. The power to determine the nature of sexual relationships held by an employed man over an unemployed destitute women is clearly greater than that held by a business executive over a lower paid colleague. The power relationship between a man and a woman with a difference of R1000 in income (in favour of the man) would be very different if the women earned nothing or if she earned R10000. The lower levels of power reduce the risk of relatively wealthier men relative to relatively poor men, as the difference in number of partners between men is likely to be smaller. This reduced directly the gradient of the relation between relative income and risk for men and indirectly for women through the reduction in partner risk. The positive slope is likely to remain as a result of the attractiveness factor, that is women's risk being determined by the relative income of their partners: on average the higher the women's status, the higher the status of her partner and therefore the higher her risk. The relationship would, therefore, be better represented by the following:

Figure 2.7 Decreasing role of relative income

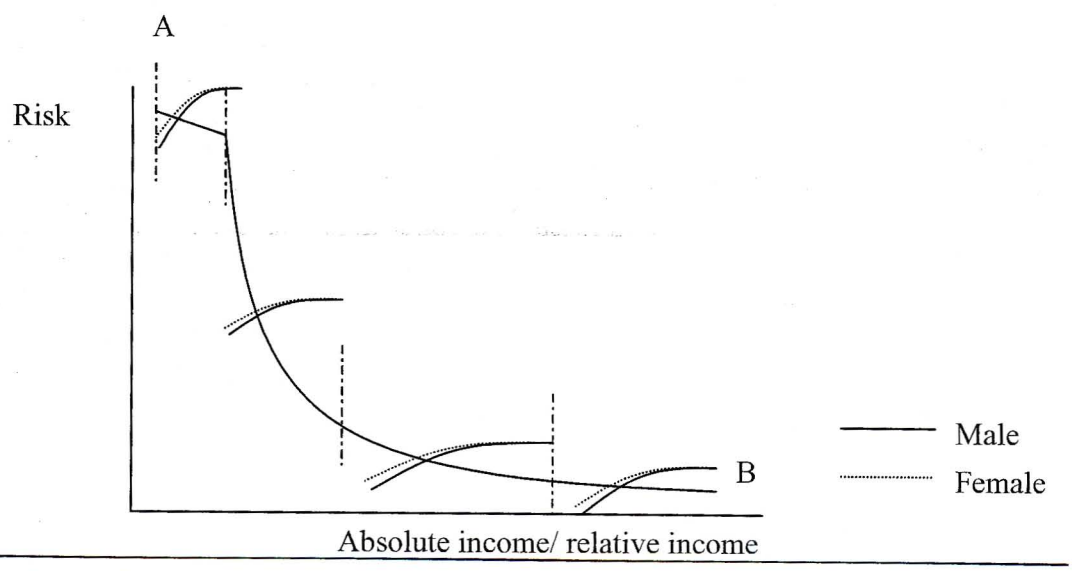

in the following discussion.

This representation shows the increase in risk associated with an increase in relative income, decreasing as absolute income increases. This representation is still oversimplified. It shows the entire society divided into neat and exclusive reference groups. Sexual networks can, however, link many individuals across income levels, countries and even continents. The incidence of HIV across the world bears testimony to this. This is not to suggest that an individual's reference group is the entire world, but rather that reference groups can overlap; therefore, while two individuals may be in different reference groups they could be linked by a common overlapping group. An individual's reference group, as defined in the introduction to this chapter, is the group of people from which an individual selects their sexual partners, plus the group from whom their prospective partners select their partners. The overlaps mean, aside from linking groups together, that two individuals may be at similar levels of absolute income but, due to the differing characteristics of their reference group, one might be relatively deprived while the other is relatively wealthy. The figure can be redrawn to include these overlaps, showing how individuals can have the same absolute income but different relative incomes and how, via many groups, the bottom reference group is linked to the top group.

Figure 2.8 Absolute and relative income with overlapping reference groups

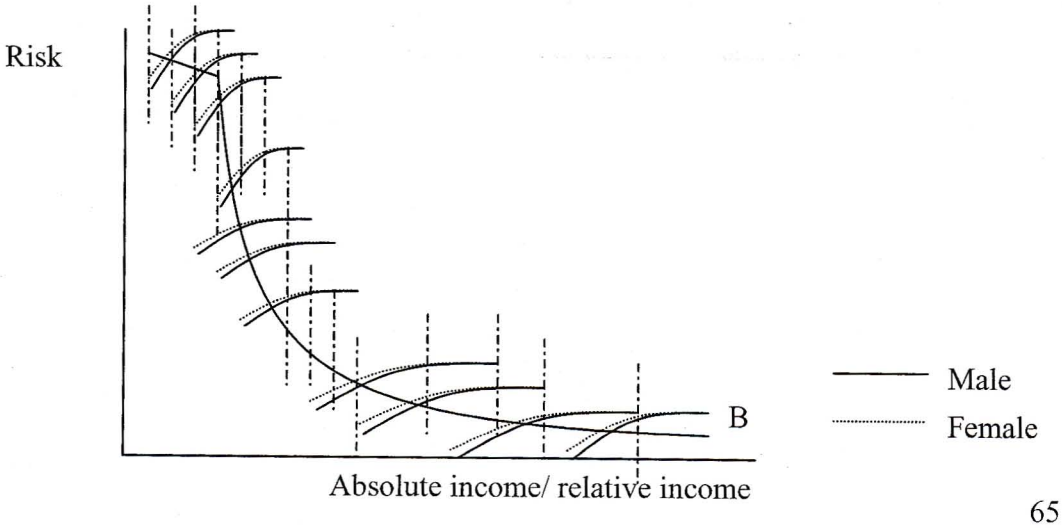

The above figure presents a more realistic representation of risk of HIV infection in relation to absolute and relative income variations. Multiple reference groups exist and overlap across the range of absolute income. For any given absolute income, therefore, there will be many individuals who are members of different reference groups, with different levels of average absolute income and therefore different relative income levels, despite having the same absolute income. This, of course, means that at any income level there will be a number of different individuals with different levels of risk. The following section speculates what will be observed if income levels are related to prevalence of HIV.

2.5 Observed relationship between income and HIV risk

The previous figure displayed the relationships between absolute and relative income and risk of HIV infection. If a cross-section of a population were taken and income and HIV status were recorded, and the above figure redrawn with income on the horizontal axis and HIV prevalence on the vertical axis, the observed relationship would be somewhat different. Unlike in many of the previous graphs, there would be only one observation for each level of income. An individual is either infected or not, their probability of infection cannot be observed. To approximate the probability, the prevalence of individuals at each income level can be recorded, effectively estimating the average risk at each income level. This is estimating the average between groups as there can only be one prevalence value for each income level.

If the relationships between income and risk described in this chapter are correct, at very low levels of income individuals at the bottom of the income range would have lower prevalence than those with slightly higher incomes. At very low levels of income the protective effect of absolute income, as previously discussed, is weak, but within groups the relative income effect is strongest at low levels of income. While individuals are

drawn from different groups in a population within all of those groups at low levels of income, as income rises the dominant influence is likely to be that associated with relative income. At higher levels of income, the relative income effect will be difficult to identify, as an increase in income will mean an increase in relative income for some and a jump in reference group for others. As income increases, the protective effects of increased absolute income are realised, quickly at first and, as diminishing marginal returns set in, more slowly. Reference groups will overlap and at any point along the income scale there may be individuals with the same absolute income, but different relative incomes and thus different levels of risk. These overlaps flatten out the increasingly weak effect of relative income variations.

The predicted observable relationship described above is displayed graphically in figure 2.9:

Figure 2.9 Predicted observable relationship between HIV prevalence and income

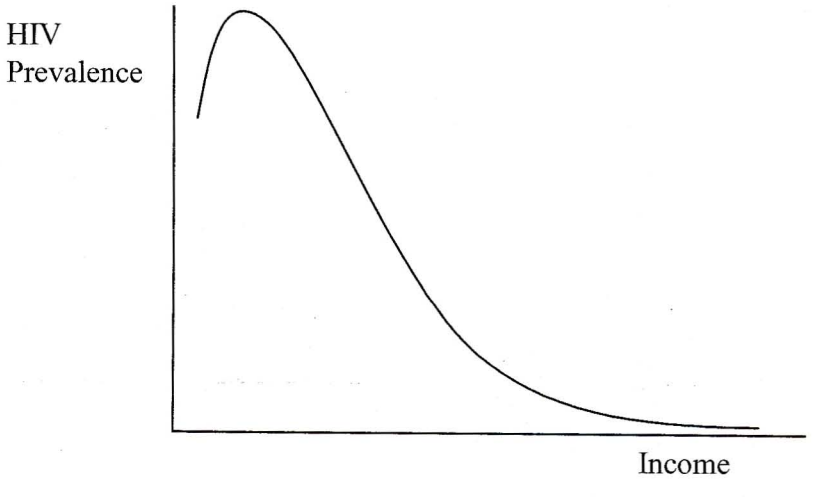

Arguably, HIV prevalence should increase at the highest levels of income, as these individuals must also be the relatively wealthy of their reference groups. Two issues should, however, be noted: firstly, risk and therefore prevalence at that point are likely already to be so low that even a doubling in risk would be barely noticeable. Secondly, it

depends on the population under examination. For example, if the population were that of a country, it would be expected that the elite would select sexual partners not only from within the country but also from other countries. Their reference group, therefore, exceeds the boundaries of their country and they may be relatively wealthy or deprived members of this group. If, however, the population were that of the world then indeed the extremely wealthy would be expected to have higher risk than other members of their reference group, albeit from a small base.

2.7 Summary and hypotheses

The probability that an individual is infected with HIV is determined by the number of sexual contacts they have with an HIV positive partner, multiplied by the probability that they will be infected by each contact. Income affects both of these variables in a number of different ways. In previous work, researchers have found both positive and negative relationships between HIV infection and income. This chapter has attempted to build on the health economic concepts of relative and absolute income, reference groups, and neighbourhood influences on behaviour and to apply these to the problem of explaining the role of income in the spread of HIV.

Absolute income is an individual's or group's income independent of others. Relative income is an individual's level of income in relation to other members of their reference group. An individual's reference group in this analysis is defined as the group of people from which an individual selects sexual partners, plus the group from which their potential sexual partners select their sexual partners.

Having made this separation, the roles of absolute and relative income can be identified. In theory, absolute income offers protection against infection with HIV. This protection is a result of a number of factors, mainly access to health care and information. Increased access to health care provides a reduced risk of HIV per contact with an HIV positive sexual partner. This protection comes primarily through reduced rates of STIs, because this is the main route through which access to health care and risk of HIV infection are

linked, though the relationship is not linear. At very low levels of absolute income and associated health care, increases in access to health care may not be of an appropriate nature to reduce risk; consequently, an increase in access will reduce risk very little. At the point where STI treatment becomes available, increases in access lead to large decreases in risk; however, once this treatment is available further access to health care offers little protection.

The second pathway through which increases in absolute income offer protection is access and ability to interpret information. Higher income leads to access to more sources of information and is also correlated with higher rates of education, which leads to better interpretation. Information about HIV can be used to change behaviour so as to protect against infection. Again the relationship is unlikely to be linear, as there are certain changes that an individual can make in their behaviour to reduce risk after which additional information offers little additional protection.

Health seeking behaviour, availability of health services, and sexual behaviour are shaped largely by an individual's community. It has been argued, therefore, that the protective influence of increases in absolute income is the result of increased group average income rather than increases in individual income. Based on this assertion, absolute income is being dealt with as a group rather than an individual variable, it is argued that the average absolute income of a group determines their average risk.

In contrast to absolute income, which is argued to determine the average risk of a group, relative income is argued to determine the distribution of individual risk about that average. Moreover higher relative income, as opposed to absolute income, is argued to lead to greater risk of infection for individuals relative to individuals in the same group but with lower income. This increase in risk comes about through linked but slightly different routes for men and women. For men, those with income relatively higher than others within their reference group are at greater risk of infection because their income increases their power over women in their reference group and increases their attractiveness relative to other members of the reference group. Both these factors

influence the number of contacts with HIV and therefore risk. For women, however, those with higher relative income are at greater risk than those women, in the same reference group, with lower income, but this higher risk occurs for different reasons. According to the principle of positive assortative matching, a woman is likely to partner a male of similar or higher socio-economic status to themselves. So, as a women's socio-economic status increases, the average status of her partner increases making him a higher risk partner, therefore increasing her risk of infection. Relatively deprived women are also at risk associated with the power men have over them. This is, however, unlikely to be the dominant effect, as most women are infected by their partners. Other things being equal, the power of men over women in groups with higher absolute income is less than the power of men within groups with lower absolute income; the relationship between risk of infection and increased relative income is, therefore, weaker.

When considered in conjunction with one another, the roles of absolute and relative income in determining risk can be represented graphically as in figure 2.10a (a reproduction of 2.8). In a cross-section of a population the observed relationship between income and HIV prevalence would be similar to figure 2.10b (a reproduction of 2.9).

Figure 2.10a Absolute and relative income and individual risk of infection

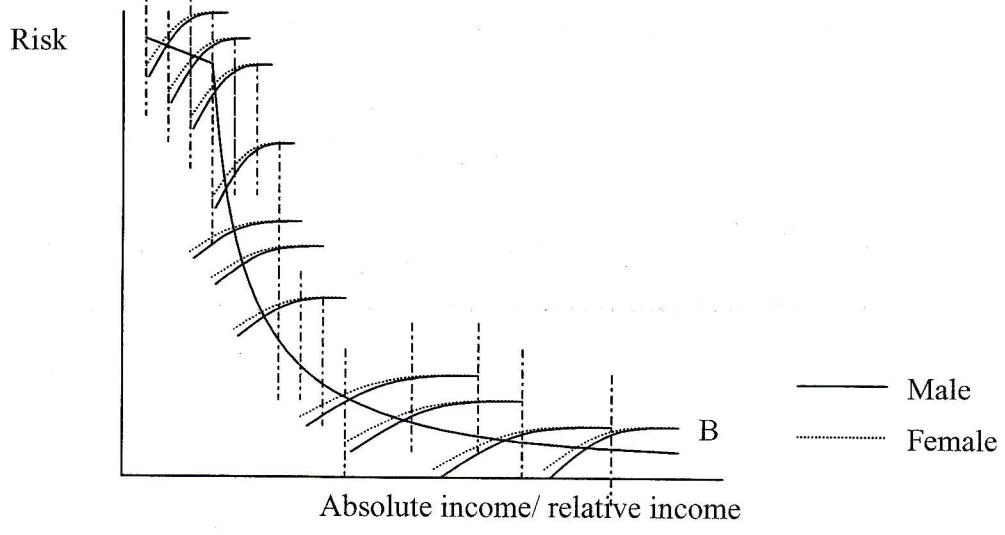

Figure 2.10b Predicted observable relationship between HIV prevalence and income

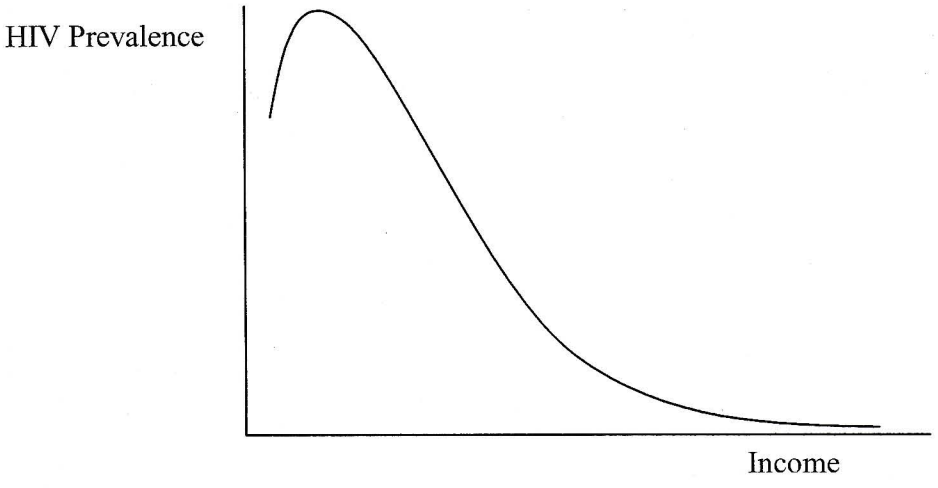

The above relationships can be represented mathematically as follows:

Theoretical relationship

Prob(Individual HIV) = contacts * probability (10)

Contacts = f(AY, rY, X) (11)

Probability = f(AY, rY, X) (12)

Prob(Individual HIV) = f(AY, rY, II, SII, X) (13)

where Prob(Individual HIV) $\propto$ 1/AY, rY, II, SII (14)

Observable relationship

HIV prevalence = $\beta_0+\beta_1 Y-\beta_2 Y^2+\beta_3 Y^3+\beta_4 II+ \beta_5 SII+\beta` X$ (15)

Where

Prob(Individual HIV) is the probability that an individual will become HIV positive

Contacts is the number of times an individual comes in to contact with HIV

Probability is the probability of infection per contact

AY is absolute income

rY is relative income

X is some group of variables related to the spread of HIV but independent of income, religion and male circumcision for example

Y is income

II is income inequality

SII is income inequality between the sexes.

The theoretical relationship suggests that an individual's risk of HIV is determined by their relative income, the absolute income of their group, the level of income inequality and the income inequality between the sexes, plus a number of socio-cultural factors independent of income.

If the theoretical relationship outlined above exists and a cross section of a population were taken and individuals were tested for HIV, the HIV prevalence would be related to income as shown in (15). Increases in income would result in increases in HIV prevalence at lower levels of income ($+\beta_1 Y$) followed by rapid decreases ($-\beta_2 Y^2$) which become less rapid as income increases ($+\beta_3 Y^3$).

The model and theory developed in this chapter build on earlier work in this field. The discussion, however, goes further than previous work and examines how the various pathways through which income can influence risk of HIV infection fit together. The theory suggests possible reasons as to how both positive and negative relationships between income and risk could have been observed and examines how the relationships are likely to change with time. The explanation of how these different relationships could have been observed, provided by the World Bank (1997), is not rebutted by this work. What this chapter does is provide more detail to the table.

Based on the model and the discussions in this chapter, the following hypotheses can be made. If the theory outlined in this chapter is correct, these hypotheses will hold true in a mature epidemic where the primary mode of transmission is heterosexual sex:

H_1: The higher the average income of an individual's reference group the lower is the probability that the individual will be HIV positive.

H_2: Individuals in the same reference group with higher relative income have greater risk of being infected than individuals with relatively low income.

H_3: Within a reference group, women have higher HIV prevalence than men.

H_4: In a population cross section, HIV prevalence is a cubic function of income.

In the following chapter the methods that will be used to test the above hypotheses will be discussed and outlined and the data described. The results of the application of these methods are presented in chapter 4, in chapter 5 the theory developed in this chapter is expanded, and the implications of the results are discussed in the final chapter of this thesis.

Chapter 3

Method and Data

3.1 Introduction

The previous chapter used the health economic concepts of absolute and relative income, reference groups, and neighbourhood influence on health to explain the relationship between the income of groups and individuals in determining their susceptibility to infection. Based on this theory, a number of hypotheses were identified. These hypotheses are tested in the following chapter. This chapter outlines the methods that were used to test these hypotheses and ultimately the theory that relative and absolute incomes influence an individual's risk of HIV infection in very different ways.

The analysis of the data for their relevance to the theory being examined can be broken down into three parts: review of existing studies, antenatal clinic data, and company surveys. Existing studies, although not designed to test these hypotheses, provide useful information. Antenatal Clinic Data are used to examine relative income effects and the Company Data to investigate absolute income effects. No primary data were collected for this analysis. The data sets used are by no means ideal, but they were the best available in South Africa at the time the analyses were conducted. To better understand their limitations, a section is included in this chapter which describes an ideal, while still realistic, data set. The following sections outline how each part of the analysis was conducted.

3.2 Review of existing studies

A number of studies relating income and socio-economic status to HIV infection have been conducted. Many of these were discussed in chapter 1. The following chapter presents a review of two examples of these studies in relation to the theory outlined in chapter 2. The hypotheses postulated in the previous chapter are compared to the results

of these studies and evidence of contradictions or reinforcements of the theory are identified and discussed.

Firstly, the results of an example of an individual level analysis of the risk factors associated with HIV infection, including socio-economic status, are compared to the predictions of the theory under examination here. The Ainsworth, et al (1998) study was selected as the example of an individual study. This study was selected, because it included an analysis of socio-economic status and risk for both men and women.

Secondly, an example of a study examining the links between group variables, including income, and HIV prevalence was selected and the results compared to the predictions of the theory outlined in the previous chapter. The Over (1998) study was selected as it represents the most in-depth study of its kind available.

The reason for selecting one individual study and one group study was because they appear to contradict one another. The Theory of the Dual Role of Income in the Spread of HIV attempts to provide a detailed explanation for this apparent contradiction.

These studies are, however, not designed to test this theory. Furthermore, many of the individual level studies are old and their relevance to current situations is questionable. More recent data are therefore necessary to test the hypotheses.

3.3 An ideal data set

In addition to existing studies, two other sources of data were used in the analysis: data from antenatal clinics and company survey data. Both of these data sets will be described and the methods applied to them discussed. They both, however, suffer a number of limitations. To understand these limitations better, it is useful first to outline what data and method would best be able to test the theory. The available data sets can then be compared to this reference point, clearly showing the strengths and weaknesses in it. This section describes the characteristics of a realistic, ideal data set. The following

sections describe the available data sets and the methods that will be used to analyse them and compare them to the ideal set, thereby identifying the limitations.

The nature of the theory requires the simultaneous examination of individual and group variables. A similar problem exists for studies examining the neighbourhood impact on health. Determining whether poor health in poor neighbourhoods is contextual or compositional cannot be done with the use of aggregate data at the neighbourhood level (Ross, 2000). It is necessary to examine individual characteristics, but at the same time place them in the context of the neighbourhood in which individuals live. If individual data alone are used, all that is being shown is the impact of individual socio-economic variables on individual health, or worse it will show the combined effect of individual and neighbourhood effects but associate both with individual characteristics. If aggregate data are used, the impact on health of living in a poor neighbourhood is indistinguishable from the impact on health of the low average income of the neighbourhood residents. For this reason, investigations of this nature require multilevel data where the unit of analysis is the individual but, in addition to individual characteristics, data on the characteristics of the individual's neighbourhood are linked to the individual observation, so each observation contains data on the individual and their neighbourhood (Ross, 2000).

An ideal data set for this study would, therefore, need to contain information on individual characteristics and characteristics of the individual's reference group, for use as explanatory variables. At the individual level the necessary characteristics are easily identified and uncomplicated. Basic demographic information is required:

- **Age** – because HIV prevalence varies greatly across age groups, an issue which will be discussed in relation to the theory under examination in the results section.
- **Sex** - because it is necessary to differentiate between sexes to examine whether the impact of income and sex-based inequality are as the theory suggests.
- In some contexts, **ethnic group** would also be required. In the South African context, on which this thesis focuses, there are many reasons why HIV prevalence differs between ethnic groups. Some of these are independent of income while others are

not. In the analysis it is, therefore, necessary to differentiate between the groups to avoid omitted variables biasing the results.

Other than demographic variables, a number of individual socio-economic characteristic variables would also be ideal. The most obvious is individual income. Education levels are also important. The theory suggests that increases in absolute income are protective, partly because of the correlation with education. It is, therefore, useful to differentiate between the two influences, to determine how much protection is derived from income independent of education. Other ideal socio-economic variables at the individual level relate to employment status, sector of employment, and whether an individual is a migrant or not. Different types of employment may affect individual risk, for example truck drivers are at higher risk as a result of behaviours associated with their employment. The underlying reasons for this, as with ethnic groups, relate to factors correlated with, but with no causal link to, income. Controls are therefore required.

An ideal data set would also include information on individual sexual behaviour as part of an effort to identify risk factors associated with various levels of income and education. Furthermore, the information can be used to investigate the role of the reference group in determining individual behaviour, to identify behaviours associated with the group's average income rather than the individual's: for example, investigating individuals with the same level of income but different behaviours, resulting possibly from being members of different reference groups.

The final individual variables would be HIV status, as the dependent variable, and STI history. These variables are an obvious necessity. Although it is theoretically possible to estimate individual risk of infection based on knowledge of sexual behaviour and group HIV prevalence, it is difficult; the data requirements are large and the margin for error great. Therefore, in an ideal data set, it would be best to collect individual HIV status.

The above individual information would need to be linked to information about the individual's reference group, as further dependent variables. This would include the

individual's reference group's average income, which is necessary for measuring the individual's position relative to others in order to estimate relative income influences and also as a measure of absolute income. This information should be further supported with the average education level of the reference group. A measure of male/female inequality is also needed as part of the relative income investigation. Migrant worker levels, both in terms of in and out migration, are essential in identifying high-risk groups and isolating the impact.

Based on the above data set, regression analysis could be conducted to isolate the roles of absolute and relative income. Hypotheses on income inequality could also be tested, as well as possible differences between the HIV prevalence of men and women. Cross-sectional analysis could be used to examine the combined role of absolute and relative income on HIV prevalence.

No such data set presently exists; even if collection were attempted many practical difficulties would present themselves. Foremost among these is the identification of individuals' reference groups, which is required as a number of reference group characteristics are necessary for the analysis. While other data sets exist which contain much of the rest of the data required, such as individual data on income, education and employment status, they are not linked to HIV infection data. Effectively, the dependent variables and the explanatory variables have no way of being associated with one another. If a way could be found to link the existing data sets, the problem of identifying reference groups would still remain.

The problem of reference group identification in this context is serious but less so than for those investigating the relative deprivation hypothesis. For those researchers it is necessary to position an individual in terms of their income relative to those with whom they compare themselves. As Deaton (2001b) notes, individuals could be comparing themselves to members of their community or to others within their workforce or possible to those they read about in the paper or see on television. Efforts to proxy reference groups by geographical location, or age cohorts, are difficult to justify. In the

context of this study, reference groups, while still impossible to fully identify and to justify proxies for, are much simpler.

According to the definition of a reference group in this thesis, an individual's reference group comprises those individuals from which they select sexual partners plus those from which their potential sexual partners select their partners. Cyber sex does not spread HIV; there has to be the possibility of physical contact between reference group members. This, therefore, already deals with the difficulty of the possible inclusion of television stars in most reference groups as, while individuals may compare themselves to the stars, few have sex with them. Television stars and other may influence sexual behaviours, but they are not, in most cases, involved in them. The possibility, therefore, of using populations of geographical areas as proxies for individuals living in that area's reference group becomes simpler to defend. The use of a population of a geographical area as a proxy for a reference group does, however, pose two difficulties: how big a geographical area and who from a geographical area to include.

In societies with high levels of migrant labour and general mobility, the group of people from which individuals could select partners is likely to be geographically widely dispersed. An ideal data set would base its collection of reference group data on reasonable assumptions about sexual networks. These assumptions would differ according to the socio-economic position and sexual norms of the population under examination. Such assumptions will also have to be made in the following sections when identifying individuals in relation to their reference groups in the data that are available.

In deciding whom to include from a population within a defined geographical area it must be recognised that by far the majority of individuals do not view all residents of their area as potential sexual partners. Firstly, reference groups are bounded by age. Young children and the elderly are not viewed, by the vast majority of the sexual active population, as potential sexual partners. Secondly, there may be divisions based on socio-economic, ethnic, religious or cultural factors. Again, some reasonable

assumptions are required to identify an individual's reference group, or at least a reasonable proxy for it.

Unfortunately, no data with the ideal variables and good assumptions about reference group size and composition, which could be used to test all the hypotheses from chapter 2, are likely to be generated. Efforts have been made, however, to use a variety of data sources to test as many of the hypotheses as possible. The first of these is the combination of data from antenatal clinics with census data on socio-economic characteristics of regions.

3.4 Antenatal clinic data combined with the 1996 census.

3.4.1 Data

All the analysis conducted especially for this thesis uses South African data. The first data examined were from the 1999 survey of HIV prevalence among women attending antenatal clinics (Department of Health, 2000), in two neighbouring South African provinces. The provinces are not named due to the confidential nature of clinic level data. The analysis of these data aims to investigate primarily the role of relative income. Owing to a number of limitations, the data could not be used to examine the role of absolute income. An alternative data set is used for this purpose. The data are only for women and the analysis is therefore limited to the role of relative income in determining the risk of women becoming infected with HIV, although some of the group variables included are relevant to both sexes. The 1999 data set was selected for two reasons. Firstly, it was the most recent data set of this nature available to the author at the time of the analysis. Secondly, it fits well with the 1996 census data. HIV has a long incubation period so the infections recorded in 1999 would have likely occurred over the previous six to eight years. The 1996 census falls in the centre of this period. Clinic HIV prevalence in the sample varied from the low teens to the mid-forties.

Once a year, in October, the Department of Health conducts a survey of HIV prevalence among women attending state antenatal clinics. The survey was first conducted in 1990. Blood specimens are taken from women attending the clinic for the first time in their current pregnancy, on a voluntary, unlinked basis (Dorrington and Johnson, 2001). Blood is collected from these women routinely for testing for syphilis and rhesus, as part of the first antenatal clinic visit, and the HIV testing is included in this procedure for the purpose of the survey. Currently over 16 000 blood specimens are taken from women at 400 clinics across the country (Dorrington and Johnson, 2001). The clinics are selected as sentinel sites based on a National Protocol, which was revised and improved in 1995 with the help of the Medical Research Council. In addition to HIV status, the age, ethnic group and education level of the women tested are recorded, no other personal information is collected. Once tested, the HIV test results cannot be linked to individual women but only to the few personal characteristics collected and the clinic where the sample was taken.

As the survey is conducted only in state antenatal clinics, the data have a number of limitations. As a sample of the entire population they over-represent the poor and provide no information on the prevalence of wealthier women who attend private clinics (Dorrington and Johnson, 2001). As a result of South Africa's economic history, the over-representation of poor women also means an over-representation of black women.

Two provinces were selected for this analysis for two reasons. Firstly, the Department of Health in South Africa releases the results of the antenatal clinic surveys only as national and provincial averages and disaggregation is available only by age. Clinic and individual (anonymous) data have become available unofficially for only a few provinces. This limited the available selection. The first province was selected because of its size and its relatively homogenous[5] make up. The analysis is concentrated on income and not factors such as religion or circumcision; the more homogenous a sample the less likely these factors are to be determinants of variation in HIV prevalence. Data

[5] The populations of both provinces are relatively homogonous in as far their populations share largely a common language and history.

for a neighbouring province were available and were included to increase the variance in the sample. Although this increases the variance of social-cultural variables, it is hoped that this is minimal, given that the provinces are neighbours and circumcision is not common in either.

The antenatal clinic data, therefore, consist of the HIV status of women attending the clinics plus their age and level of education, but unfortunately not their income. Each individual record is linked to the clinic from where the sample was taken.

Despite its limitations, the survey provides the largest source of HIV data in the country. Unfortunately, there is very little individual information. It would have been very useful to have women's employment status, their income and the income of their partners'. Since education and income are highly correlated, for the purpose of this analysis education is used as a proxy for the individual's income. It has been discussed earlier how education and income have collective as well as individual implications of risk. The implications of including only education are discussed later in this section. The data therefore have at least a proxy for individual income, unfortunately however, there is no reference group information.

Similar to the problem faced in estimating the neighbourhood influence on health, it is necessary to link individual data to variables at the reference group level. Before a link can be made, the individual's reference group has to be identified. A reasonable assumption could be made that individuals' reference groups were made up mainly of people in their community. The antenatal clinic data should therefore be linked to local variables. This, however, ignores the high levels of migration and general mobility, and the sexual networks that result; migrants' reference groups are far larger than the number of people in their local community. It could, therefore, also be reasonably assumed that references extend across the entire sample. This would imply a need to stratify data for the whole geographic area and somehow link individuals to one of the strata as their reference group. This option is, however, very unlikely both practically and theoretically for the following reasons.

The theory developed in the previous chapter argued that differences in prevalence between groups resulting from absolute income occur because of differences in health care access and behaviour based on available information. The entire sample is based on women attending state antenatal clinics; they therefore have similar access to health services as well as similar health seeking behaviours. The sample is of a relatively homogenous population, as attendees at the clinic are of similar socio-economic status or at least are accessing health services of a similar quality, and similar ethnic background. Variations in risk are, therefore, unlikely to be a result in variations in absolute income. Placing women in socio-economic strata would therefore be pointless, because even if they could be allocated they would almost all be placed in the same stratum. This is the same reason why it is acceptable to use the individual's education level. If individuals are drawn primarily from the same reference group, or at least groups with similar average absolute income, then they have a similar level of group average risk associated with the current level of average absolute income. Higher rates of individual education are, therefore, unlikely to offer additional protection as they still have the same access to health services and probably face a similar set of socially acceptable behaviours. Effectively this means that, if education and income are correlated, then as education increases, the effects of an increase in relative income will be observed. In South Africa education has been found to be correlated with income, although it is important to control for age and marital status, which also play a deterministic role in relation to income (Case and Yogo, 1999). For women, the return on a year of education was found to be, on average, 14 percent of income (Case and Yogo, 1999). In this data set age can be controlled for but not marital status. Marital status affects income primarily through child bearing. Given that all the women in the sample are pregnant, this variable is unlikely to account for significant variation and omitting it is unlikely to make any significant difference.

Practically identifying break points for strata would have been difficult. As all the women have similar access to health care the breaks would have to be made based on behaviour changes relating to income levels of reference groups. As no behavioural data

associated with income were available for the region, separating the sample into different reference groups with different average absolute incomes would have been impossible.

Although there are theoretical and practical difficulties associated with identifying reference groups in order to investigate the implications of relative income on risk, the women in the sample have to be placed in relation to their reference group. To this end some assumptions as to the appropriate reference group for each individual must be made. As an initial step it was assumed that the women were from either the same reference group or at least they were from different groups with similar absolute income. The analysis was, therefore, conducted twice; firstly, assuming that the local community, in this case represented by the magisterial district in which the clinic the women are attending is located, was the appropriate reference group. Then using education as a proxy for individual income, women were placed in relation to others attending the same clinic by measuring their education as a percentage of the clinic average. The second analysis included women's education directly, assuming that the reference group spanned the entire sample, thus placing the women's education level relative to the average education of all the women in the sample.

It is unfortunate that no measure of partners' income or education was available. In chapter two much was made of how the partner's income, for women, was an important determinant of their risk. However, as the hypothesis on the relationship between relative income and risk for women already includes the risk associated with the partner it can still be tested with these data. It is an attempt to examine the total risk of women in relation to their relative income including the risk resulting from their partner's income and associated behaviour.

Whichever assumption regarding the size of the reference group is made, it still means that the data set contains only individual level data and it is necessary to combine group level data with each data set. Even if the entire sample represents one reference group, these local conditions (group data) will still influence the risk of HIV infection for individuals in these areas.

The group variables are based on data collected in the 1996 South African Census conducted by Statistics South Africa. The census was used for two reasons; firstly, because it was conducted at the mid-point of infections, that is the majority of infection in 1999 would have occurred between 1993 and 1996, given an average incubation period of 8 years and a decline in fertility as the virus progresses. The second reason for using the census was its sample size. The data for the group variables were necessary at a level at which they could be linked to individual clinics. Other national surveys, such as the October Household Survey, do not have sufficient sample size to disaggregate to this level. The 10 per cent sample, produced by Statistics South Africa, was used to increase the speed of calculations. To estimate group variables representative of the reference group for women attending each clinic would have required accurate information on the catchment area for each clinic. This information was not available and a simplifying assumption was necessary. For the purpose of this analysis, data from the magisterial district in which each clinic is located were used as the closest available data on the local community of attendees; fortunately, there were no cases of more than one clinic in a magisterial district. Based on the common variable, magisterial district, the antenatal clinic data and the census variables could be combined into a single data set.

The data for the magisterial district were refined in a number of ways to better represent individuals' reference groups. Firstly, averages of the appropriate variables included only adults in the sexually active age group, this was assumed to be 15 – 49 years of age. These were further reduced to include only black Africans, as by far the majority of clinic attendees fall into this ethnic group. The only exceptions to this were the variables measuring total household expenditure. These were calculated for the entire population of the magisterial district. This was done because it is the wealth of the area and not that of a sub-population that is likely to play a role in attracting migrants.

A number of different variables were examined[6], including measures of sexual inequality, female poverty, population mobility and regional wealth. A list of these variables and how they were constructed is contained in the appendix.

The variables included in the final model consisted of HIV status as the dependent variable, a series of age dummies, a measure of female status, household expenditure, a measure of mobility and the education level of the women in years of education. The measure of female status was constructed as female average income divided by male average income and the result multiplied by female average income. The variable was designed to capture both the effects of inequality and income. Areas with low average female income and high inequality result in the lowest values for this variable, while those areas with low inequality and high average income produce the highest values. The mobility measure was constructed as the proportion of men living in the magisterial district who work in another magisterial district, but within the same province. A more detailed discussion of the variables is contained in the next chapter. The variables are described in the following table.

Table 3.1 Antenatal clinic data descriptive statistics

Variable	Mean	Std. Dev	Min	Max	Obs.
HIV status	0.31	0.46	0	1	4615
Age 15-19	0.18	0.38	0	1	4615
Age 20-24	0.31	0.46	0	1	4615
Age 25-29	0.25	0.44	0	1	4615
Age 30-34	0.17	0.37	0	1	4615
Age 35-39	0.07	0.26	0	1	4615
Age 40-	0.02	0.13	0	1	4615
Years of Education	8.54	3.29	0	13	4603
Years of Education	83.68	46.10	0	169	4603

[6] The basis for variables inclusion or rejection is outlined later in this chapter

squared					
Female status	135.61	106.25	9	407	4607
Male short distance work related mobility	0.03	0.05	0	0.24	4615
Monthly household expenditure	2103.48	987.16	807	4573	4615

3.4.2 Analysis

Having constructed the data set containing both individual and reference group variables, regressions examining the nature, if any, of the relationship between HIV infection and the individual and group variables were conducted. Two approaches to conducting this regression analysis were considered; the first was to use the same approach as Over (1998) and the second to use individuals' status as the dependent variable. An explanation of each option and the basis for the decision between them follows.

Over (1998) examined the variation in HIV prevalence, the percentage of the population infected at a given moment in time, between countries. He used HIV prevalence data from the US Bureau of the Census as the dependent variable and a number of socio-economic variables, including GNP per capita, as explanatory variables. Over then used Linear Multiple Regression to test his various hypotheses. HIV prevalence over different stages of an epidemic is, however, a non-linear variable and it was necessary to transform the HIV prevalence rate. The HIV prevalence rates were divided into urban high-risk and low-risk groups. For each group, a ceiling was set at a rate higher than any observed rate in that group. For the high-risk group, the ceiling was set at 90 per cent and for the low-risk at 40 per cent, these values being just above any recorded values for each group. The prevalence rates were then transformed as follows:

$Log_e [y/(c-y)]$

Where:
Y is the infection rate
C is the ceiling

Rates of 0 were changed to 0.1 to prevent those observations being dropped from the analysis.

The same approach could be used with the current data set. Clinic HIV prevalence could be calculated and transformed, based on some appropriate ceiling. The clinic prevalence could then be combined with the variables calculated for the magisterial district in which it is located to form a single data set linked by the common magisterial district variable.

Although such an approach would yield interesting results concerning the importance of the group (community) variables, it would do nothing to support or undermine the theory under examination. It has been argued that given the nature of the sample there will be little, if any, variation in risk as a result of differences in absolute income. What can be examined is the role of relative income, but relative income is an individual variable measuring position within a group. HIV prevalence, however, is a group measure effectively providing a proxy for average group risk, while providing no information on the distribution of the risk within the group.

The second approach of using individual data does, however, allow some of the hypotheses to be tested. This approach allows for the inclusion of both individual and reference group data. This is achieved by using the individual's HIV status as the dependent variable and the individual's age and education as explanatory variables, plus the variables calculated from the census.

An individual is either HIV positive or not; the choice of such a dichotomous dependent variable necessitated the use of models that could deal with such a variable. Three common approaches are available, the linear probability model, the probit and the logit.

Consider a model of individual HIV infection where HIV=1 if the individual is infected and 0 if they are not. The probability that they are infected is determined by some vector of variables x so that:

$$\text{Prob}(HIV=1) = F(\beta'x)$$
$$\text{Prob}(HIV=0) = 1 - F(\beta'x)$$

The challenge is to identify a suitable model for the right hand side of the equation (Greene, 1997). One option is the linear probability model:

$$F(x, \beta) = \beta'x$$

As $E[HIV \mid x] = F(x, \beta)$

$$HIV = E[HIV \mid x] + (HIV - E[HIV \mid x])$$
$$= \beta'x + \epsilon$$ (adapted from Greene, 1997)

The linear probability model, however, suffers a number of shortcomings. Some of these can be dealt with relatively easily, while others are more fundamental. The most problematic shortcoming is the possibility of a predicted value outside of the 0 to 1 range. The situation is represented graphically in figure 3.1 below:

Figure 3.1 Shortcomings of the linear probability model

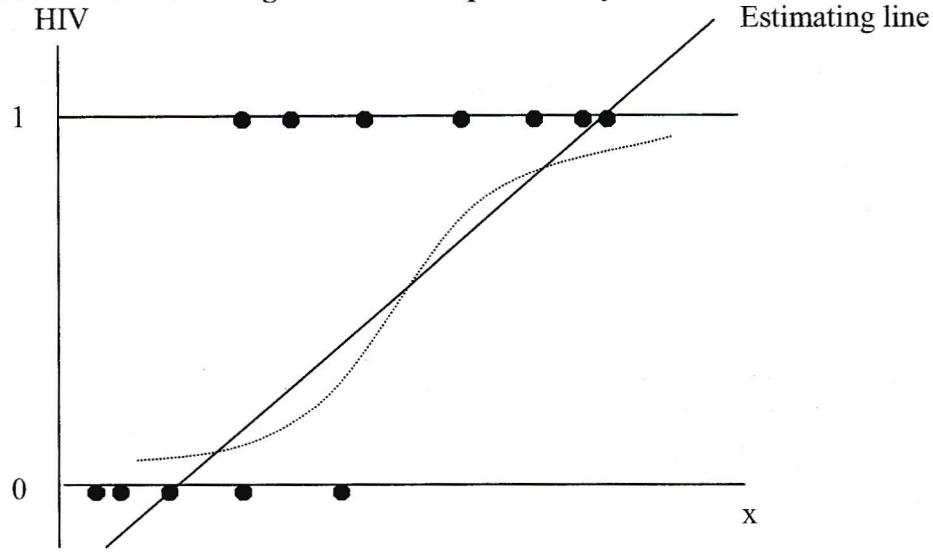

Source: adapted from (Kennedy, 1999).

At high values of x the estimating line predicts probabilities in excess of 1 and at low values of x it predicts negative probabilities. What is necessary is to estimate a function that will always predict values within the 0-1 interval, such as the dotted line in the figure above. For any regressor vector the functional form should be such that:

$$\lim_{\beta' x \to +\infty} \Pr ob(HIV = 1) = 1$$

And

$$\lim_{\beta' x \to -\infty} \Pr ob(HIV = 1) = 0$$

(Adapted from Greene, 1997)

Any proper continuous probability distribution which is defined over the real surface will meet these requirements (Greene, 1997). Two distributions are commonly used: the cumulative normal function (the probit model) and logistic distribution (the logit model). The functions are as follows:

Probit

$$\Pr ob(HIV = 1) = \int_{-\infty}^{\beta' x} \phi(t)dt$$

$$= \Phi(\beta' x).$$

Logit

$$\Pr ob(HIV = 1) = \frac{e^{\beta' x}}{1 + e^{\beta' x}}$$

$$= \Lambda(\beta' x).$$

There is little theoretical guidance to inform the selection of one model over the other. The two distributions are very similar except in the tails (Greene, 1997). For central values of β'x there is, therefore, little difference in the results of the two. The logit model generally predicts higher probabilities to y = 0 (HIV = 0) when β'x is small and smaller probabilities when it is high, when compared to the probit estimations (Greene, 1997).

For the purposes of this analysis the logit model was used. The selection was made because the logit model is slightly easier to interpret and, with no theoretical reason to guide the decision, ease of use was used as the deciding factor.

3.4.2.1 The education variable

Logit regressions are ideal for use in this context with individual HIV status as the dependent variable. The explanatory variables used can be drawn from the census, linked to the clinic that the individual attends and the relevant individual data. In this case the individual data comprise the women's age and education. The use of these variables will be discussed in detail in the results section. The use of the education variable, however, requires discussion more immediately.

In this analysis, education is used as a proxy for individual income. To test hypotheses concerning relative income it is necessary to define the individual's reference group, or at

least to place their income relative to others in their reference group. As noted previously, it can be argued that the entire sample represents one reference group; a case could also be made for assuming that the reference group is much smaller. In the first case, education could be directly included in the regression. If, however, the reference group was assumed to be best represented by the women at the same clinic, the education data would have to be transformed as follows:

TE = E/CEA

Where
TE is the transformed education value
E is the years of education
CEA is the clinic education average

The transformed variable measures education in relation to other women attending the same clinic. For values below 1 the woman is considered relatively deprived and for values above 1 relatively wealthy. The two options represent the two estimable extremes of possible assumptions of reference group size: the entire sample and the clinic alone. Both were used and the results compared. It is possible that clinic attendees at the same clinic are from different reference groups, data on which to base any division within clinics were, however not available.

3.4.3 Summary of antenatal clinic data analysis

Logit regressions were conducted using individual HIV status as the dependent variable. The explanatory variables were drawn from two sources, additional information on individuals attending the clinics and the 1996 census. The individual information consisted of the women's age and years of completed education. These individual data were linked to data drawn from the census. The link involved a number of steps. Firstly, the census data were limited to black adults between 15 and 49 years of age. This was done to better represent a woman's reference group. Secondly, the averages of a number

of variables were calculated at the magisterial district level. Thirdly, the magisterial averages were linked to individual records by linking the individual to the clinic that they attended and the clinic to the data for the magisterial district in which it was located. The regressions were estimated twice, once measuring individual education against the entire sample and once with it against the average of women attending the same clinic. This was done to deal with the two possible extreme assumptions that could be made about the appropriate size of the reference group.

Based on the above, the following hypothesis was tested:

H_2: Individuals (women) in the same reference group with higher relative income have greater risk of being infected than individuals with relatively low income.

The above analysis, although useful, fails to address a number of issues. Most notable the data and model do not examine the role of absolute income. For this reason a second data set was identified and analysed.

3.5 Company X data

3.5.1 Data

The limitations in the analysis of the antenatal clinic data stem from the narrow socio-economic band from which the individuals were drawn. They all have similar access to health services and therefore have little variation in terms of protection offered by absolute income. To measure the impact of absolute income on individuals' risk of HIV infection it is necessary to have comparable data on HIV prevalence across a wide income range. Ideally, such a data set would contain data on individuals from all income categories. No such ideal data set could be found for South Africa. What were available, however, were the results of a large company's internal HIV surveillance exercise. Although this does not contain a full range of income, as all those tested were employed,

it did give the opportunity to investigate the role of absolute income, as the company employed low paid unskilled labour through to highly paid executives. The data set also contained observations from across the country and across different ethnic groups.

The study and data from company X are described in detail in Colvin et al (2000), the following is a summary based on the Colvin article. The study was conducted to provide the company with an estimate of HIV prevalence and what the associated risk factors were. This was to inform an assessment of the economic impact on the company of the illness and death that follow HIV and to serve as a baseline to evaluate future interventions designed to prevent HIV transmission.

In 1999, the time of the study, the company concerned, employed in the region of 35 000 employees, spread across 175 sites, with some in every province. With the use of a probability proportional to size sampling methodology, 31 sites were selected from those sites housing more than 100 staff. In the smaller sites (less than 200 employees) every member of staff was sampled. In the larger sites, however, departments were randomly selected and the entire staff of each selected department was sampled to obtain samples in the region of 250. Although individual random sampling may have been better, it would have been logistically difficult and, by selecting entire departments, participation rates could still be calculated. Participation rates[7] are important in determining the representatively of the sample. It is easier to measure the rate at which available candidates refuse at a department level than at a site level.

Participation in the study was voluntary and verbal consent was obtained. The use of verbal, as opposed to written, consent was at the request of the trade unions. Each willing participant was asked to fill in a questionnaire and provide a sample of blood or saliva for HIV testing. The questions asked provide information on: the age of the employee, their sex, the geographic location of their workplace, their ethnic group, their job category, the type of accommodation they live in, their time away from home -

[7] Participation rates at antenatal clinics are typically very, high as blood is taken routinely, and not a cause for concern in that context.

measured as the time lapse between visits - STI symptoms, sexual behaviour and their exposure to various components of the company's prevention programme.

This collection of information and HIV status was part of a larger study, which aimed to estimate the economic costs of HIV infections to the company. At this stage, data on employee earnings were not collected. At the economic evaluation stage, however, income data were collected. But, unfortunately, these data were not linked to individual HIV status.

To examine the role of absolute income, some income variable or at least a proxy of income is obviously essential. To address the problem of the absence of any such variable in the individual data a link was made to the economic costing component of the study. From the costing study the average incomes by ethnic group, job band, age and sex were obtained. These averages were assigned to individual observations based on the appropriate average for each individual's ethnic group, age, job band and gender.

Once an income variable had been generated, one further modification to the data set was necessary. In the previous chapter the possible differences between male and female infection rates at the same level of income were discussed. For this reason, and because the sample was overwhelmingly male, it was decided to use only data from male participants. An alternative would have been to use a dummy variable to distinguish between male and female individuals; this, however, would have been difficult because the sample of female employees was very small.

Before an analysis could be conducted one issue had to be addressed. The theory argued that changes in absolute income should be thought of as changes in the average income of an individual's reference group. This was because of the group influence on behaviour and access to health services and information. The data set being used comprises individual records and the income is the income of the individual. It is assumed, however, that individuals in the sample are drawn from a number of different reference

groups[8] and that the higher the income of the individual the higher the likely average income of their reference group.

The sample was drawn from across the country, across ethnic groups and many different income levels. Clearly, therefore, the individuals within the sample are not from the same reference group. If, as has been assumed throughout the paper, individuals mix sexually with partners of similar socio-economic status, then as an individual's income increases, the expected value of his reference group's average income also increases. A sample of individuals from different reference groups and with different incomes is therefore likely to display the relationship between income and HIV infection attributed to changes in absolute income. The situation is similar to that described in the theory section, and graphically depicted in figure 2.9, when discussing the likely relationship between income and HIV prevalence in a cross section of the population.

The variables included in the final model consisted of HIV status as the dependent variable and age, ethnic group, provincial site location, STD symptom and housing dummies as well as income as explanatory variables. The data are described in the following table.

Table 3.2 Company data descriptive statistics

Variable	Mean	Std. Dev	Min	Max	Obs.
HIV status	0.09	0.28	0	1	4098
Aged under 25	0.06	0.24	0	1	4098
Age 25-29	0.10	0.30	0	1	4098
Age 30-34	0.12	0.33	0	1	4098

[8] Unlike the situation with the antenatal clinic data were the women were accessing similar health services and assuming that the protective influence of absolute income was, therefore, similar as defensible, the individuals here may be from many groups and while individuals may have similar income they may be from different groups with very different average absolute income and resultant protection. For this reason it was not possible to investigate relative income affects with these data.

Age 35-39	0.20	0.40	0	1	4098
Age 40-44	0.19	0.39	0	1	4098
Age 45-	0.33	0.47	0	1	4098
Mpumalanga	0.37	0.48	0	1	4098
Eastern Cape	0.05	0.22	0	1	4098
Free State	0.08	0.28	0	1	4098
KwaZulu-Natal	0.08	0.28	0	1	4098
Northern Cape	0.03	0.17	0	1	4098
North West	0.04	0.20	0	1	4098
Northern Province	0.13	0.34	0	1	4098
Western cape	0.09	0.28	0	1	4098
Gauteng	0.13	0.33	0	1	4098
Asian	0.03	0.17	0	1	4098
Coloured	0.08	0.26	0	1	4098
White	0.31	0.46	0	1	4098
African	0.58	0.49	0	1	4098
STI	0.06	0.23	0	1	4098
Single Sex hostel	0.13	0.33	0	1	4098
Construction Camp	0.13	0.34	0	1	4098
Own home	0.55	0.50	0	1	4098
Other accommodation 1	0.17	0.37	0	1	4098
Other accommodation 2	0.02	0.16	0	1	4098
Income	68439.45	28705.80	36492	275244	4098

3.5.2 Analysis

The analysis concentrated on identifying the role of absolute income changes. A similar regression to that conducted on the antenatal clinic data was undertaken. Individual HIV status was used as the dependent variable (infected or not). Again, as with the antenatal

clinic data, such a binary dependent variable necessitated the use of an appropriate regression model. For the same reasons as before, the logit model was selected.

The explanatory variables comprised a number of variables obtained from the questionnaires, including the individual's age and ethnic group, location of site where the individual works, the type of accommodation in which they live, and whether they had had symptoms of an STI. The dependent variables also included an estimate of the individual's income as described above.

3.5.3 Summary of Company X data analysis

The antenatal clinic data provided an opportunity to examine the influence of relative income on the risk of infection. Due to the socio-economic characteristics of the individuals contained in the sample, it was not possible to use the data to study absolute income. An example of the narrow band is that all the women in the sample had access to similar health services, in this case state antenatal care. The data from company X comprised individuals from across the country and from different reference groups, with large variation in socio-economic status. Although individual HIV status was linked to other useful individual data, the original data set did not have information on individual income.

The lack of income data was dealt with by linking individual records to income figures based on the average for their sex, age, job band and ethnic group. Assuming that individuals were not all from the same reference group and that the higher their individual income the higher the expected value of their reference groups average income, allowed for the analysis to focus on absolute income variation. Some individuals in the sample may well be from the same reference group and those with higher income would therefore be expected to be at greater risk i.e. the relative income effect. Over the range of income, however, a number of reference groups will be passed through and the average risk at continually higher levels of income would be expected to fall, i.e. the absolute income effect.

Regression analysis was conducted using logit regressions. The dependent variable was the individual's HIV status and the explanatory comprised data from the questionnaire and the linked income data.

Based on the above, the following hypothesis was tested:

H_1: The higher the average income of an individual's (in this case males) reference group the lower is the probability that the individual will be HIV positive.

It was also possible to examine some of the underlying pathways linking absolute income variations to different risks of infection. These data also allowed for the partial testing of the following hypothesis:

H_4: In a population cross-section, HIV prevalence (in this case among men) is a cubic function of income.

The sample comprises a partial cross section of the population. The lower levels of income are not included in the sample as all the individuals are employed. It is at these low levels of income that it would be expected that increases in income lead to increased risk. If this part of the population is excluded, H_4 can be rewritten as follows:

H_{4b}: In a truncated population cross-section, where low level income earners are excluded, HIV prevalence is a quadratic function of income.

While these data are insufficient to investigate H_4 they can be use to test H_{4b} which is a derivative of the original hypothesis, This hypothesis can be further investigated with data from other company surveys.

3.6 Additional company data

The data from company X were the only available data with enough detail to test H_1. Other companies have, however, done similar studies. Although the entire data sets were not available, many aggregate results have been released. These data sets do not allow for analysis that controls for other risk factors, but they are worth a brief examination, if only to see if prevalence falls with income. This is different from the above, because it is effectively a univariant analysis with no controls included. For example, HIV prevalence may fall as income rises, not because of income but because those who have the higher income might be older and be less sexually active. The available data from these other company studies is, however, presented, in the next chapter, and discussed in relation to the theory under examination for interest, despite the limitations .

3.7 Summary

This chapter has outlined the methods and the data that will be used to test the theory outlined in the previous chapter. The empirical analysis in the next chapter is divided into four sections: a review of two examples of existing studies, an analysis of antenatal clinic data, detailed study of data from a company survey and a brief review of other company data.

In the literature review, a number of studies that have examined the role of income and other variables in determining individuals' risk of infection with HIV and groups' HIV prevalence were discussed. These studies provide valuable data and the hypotheses outlined in the previous chapter are compared to the results of two examples of these studies. Many of these studies appear to contradict each other regarding the role of income. It is hoped that the examination of the results based on the new theory will help explain these apparent contradictions. Collectively, the two examples of the literature discussed in the next chapter are indicative of this contradiction.

The antenatal clinic data provide an opportunity to examine the influence of relative income on individuals' risk of HIV infection. The data set consists of individual HIV status and some demographic information, individual education and census data for the sexually active black population of the magisterial district in which clinics are located. This chapter has outlined the regressions that were undertaken with these data. Logit regressions were estimated with the HIV status of the individual as the dependent variable and the demographic, education and census data as the explanatory variables. These regressions were undertaken twice; once assuming that the individual's reference group was best represented by the immediate community and another assuming that geographically the reference group covers the entire area from which the sample was drawn. The analysis allows for the testing of the hypothesis concerned with the relative income of women.

The antenatal clinic data is, however, unable to address issues relating to absolute income. For this reason data from company X were used. This chapter described the data and the method used to analyse them. The data consisted of individual HIV status from a sample of employees from company X, which operates across the country. It also included demographic information, location of the employees' workplace, whether an employee had had symptoms of an STI, their type of accommodation, their time away from home and their position within the company. The data did not include individual income. This was resolved by assigning an income level to each individual equal to the average income for employees with the same gender, age, job band and ethnic group. Hypotheses relating to changes in absolute income were then tested based on the results of logit regressions using HIV status as the dependent variable and the other variables as explanatory variables. It was assumed that individuals were drawn from a number of different reference groups and that the higher their income the higher the expected value of their reference group's absolute income. Based on this assumption, changes in income in the data can be treated as changes in absolute income.

The company X data were the most detailed data set available from a company study. Similar studies have, however, been conducted in a number of other companies.

Although the detailed data sets were not available, analysis and discussion based on what is available was undertaken.

Chapter 4

Results

4.1 Introduction

This chapter presents results in four sections, each section examining an aspect of the theory developed in the chapter 2. The methods applied in the generation of these results were discussed in the previous chapter.

Firstly, the hypotheses developed in chapter 2 are examined in relation to the results of two existing studies. One of the primary objectives of this thesis is to explain the apparent contradiction in the empirical work already done in this area and to develop a more in-depth understanding of the role of income in the spread of HIV. It is useful, therefore, to re-examine the two examples from the empirical studies referred to in the literature review to ascertain if the theory developed is able to explain the identified contradiction by building on existing explanations.

Secondly, the chapter tests the hypothesis relating to relative income with the use of a data set comprising the results of antenatal clinic HIV testing, individual demographic data, and aggregate socio-economic variables obtained from the 1996 census.

Thirdly, the results of the regression analysis conducted using company data are presented. These results are designed to test the hypothesis relating to the role of absolute income in determining the risk of HIV infection.

Finally, data from other companies are compared and the results discussed in relation to the theory and hypotheses in question. Detailed analysis of these data sets could not be undertaken because no disaggregated results were available at the time of writing.

The results of the four analyses are drawn together in the final section. They are discussed and compared, and conclusions and limitations identified.

4.2 Re-examination of existing studies

This section re-examines the results of two of studies mentioned in the literature review. The results of these studies are compared to the hypotheses generated in chapter 2. It is argued in this section that the theory presented in this thesis more fully explains the results of, and apparent contradictions between, the various studies with regard to socio-economic status and HIV infection, by building on existing work.

It is not possible here to re-examine all the studies on the links between socio-economic status and HIV. Two studies were therefore selected as examples: the Ainsworth and Semali (1998) study of individual infections in Kagera, Tanzania and Over's (1998) cross country investigation. The reasons for the selection of these two as examples are described in the previous chapter, essentially they where selected as they provide a prime example of the apparent contradiction between group and individual studies relating HIV infections to income. The results of these studies are compared to the hypotheses posed in chapter 2 and the explanatory power of the theory is compared to work already done in the area that attempts to explain the apparent contradiction between individual and country level studies.

4.2.1 Individual risk of infection

The literature review mentioned that a number of studies have examined the link between individual characteristics and risk of infection. This re-examination concentrates on one of these studies, the Ainsworth and Semali (1998) paper. The study utilised data collected from households in the Kagera region, Tanzania, between 1991 and 1994. It compared the influence of age, education, sex, wealth and occupation on the probability of dying from AIDS or dying from another cause. The region at the time was home to over a million people; more than 80 percent of the population were living in rural areas, with the

majority income generated from agricultural production. From the information provided there is little suggestion that, it would not be safe to assume that, with the possible exception of a small elite, the protective influence of absolute income varied very little across the sample. There was no mention of variations in access to health care and by far the majority of the sample appears to be drawn from rural farming households. If it was the case that absolute income did not vary enough to account for different levels of risk it would be more likely that the influence of relative income would be observed, as it is reasonable to assume that most individuals are part of the same reference group, or different references groups with a very similar levels of absolute income.

Of primary interest from this study is the predicted probability, by education level, of dying of AIDS. Wealth and the probability of dying are not particularly relevant, as illness and deaths resulting from AIDS represent a major negative economic impact complicating the interpretation. Household wealth at the time of death is likely to be very different, in absolute and relative terms, than it was at the time of infection as households suffer numerous economic impacts during the illness and on the death of a household member.

Within the sample the relationship between education and the probability of dying observed in the study was as follows:

Figure 4.1 Male probability of dying

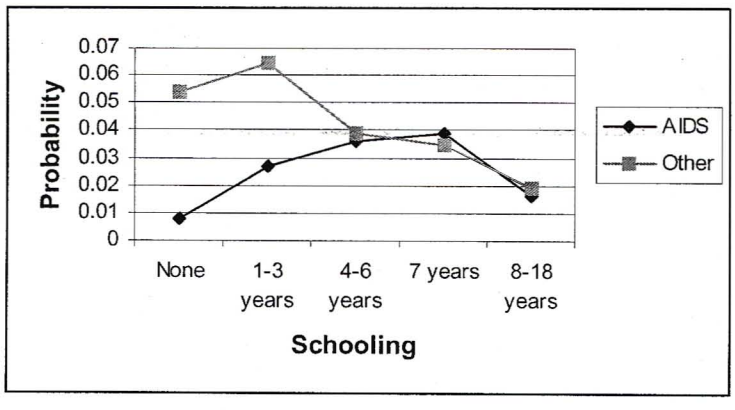

Source: Ainsworth et al, 1998. p103.

Figure 4.2 Female probability of dying

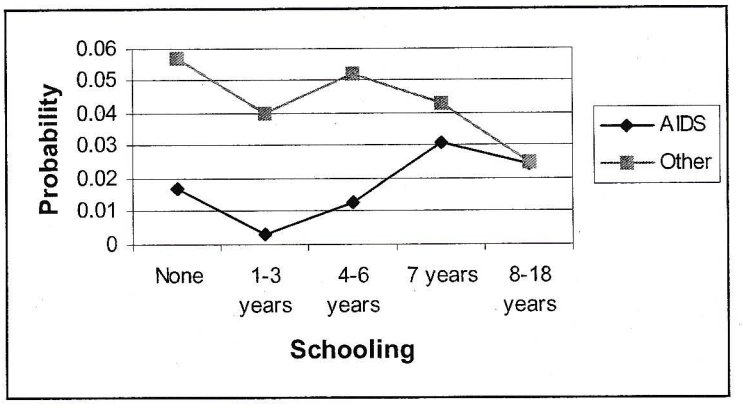

Source: Ainsworth et al, 1998. p103.

Figure 4.1 displays the relationship between the probability of dying from AIDS or from other causes for men, by level of schooling. For other causes, the probability of dying generally decreases as education increases. This is the normal expected relationship between mortality and education. The distribution of AIDS deaths and the resultant probabilities of dying are somewhat different. The probability rises up until 7 years of schooling and then falls for the group 8-18 years. This is consistent with expectations based on the theory developed in this thesis. The initial rise is in line with the argument that increases in relative income (proxied here by education) exacerbate risk of infection. The final decline in the probability of dying is also in line with expectations. Although it may be reasonable to assume that the sample reflects mainly one reference group or many reference groups with similar absolute income, it is not reasonable to assume that everyone in the sample is in a reference group with such an average absolute income. At the upper end of the education scale, the average income of individuals and resultant access to health services and information might be sufficiently different from those with lower income as to offer them increased protection as a result of their higher absolute income. The risk of these individuals would be expected to be lower because of the protective influence of absolute income. This would reduce the average risk of the individuals, in this case the males, in the sample.

Figure 4.2 displays the same relationships for women. Again, for other causes of death, there is a downward trend in the probability of dying as education increases. For the probability of dying from AIDS, the relationship is somewhat different from both the other causes of death and the male probability distribution. The theory developed in chapter 2 of this thesis suggests that female risk of infection is related to male risk. If this were the case then women's risk of dying would increase with education and then decrease, as was observed in the male sample. This was the case for women in the sample, with the exception of women with no education who had a higher probability of dying than women with a little education. Although not entirely consistent with predictions, the result is not necessarily in contradiction to the theory presented in this thesis. It may be the case that the power held by men over very poor women, such as women with no education, places them at more risk than the increased sexual partnering of the men partnered with women with slightly higher income. This would suggest that this relationship would be observed in reference groups where some women are very vulnerable, while in reference groups with higher absolute income there would be a simple upward trend in risk. Furthermore, poor women may die more quickly because of their lack of access to care and low nutritional status. In the early stages of the epidemic, when deaths were still relatively rare, this may have biased the result. Overall, there is still an upward trend in risk as education of women increases, which is what the theory predicts.

Broadly speaking, the results of the Ainsworth and Semali (1998) study fit with the predictions based on the theory of the dual role of income in the spread of HIV. The positive relationships observed between education and the probability of dying of AIDS is arguably the relationship between increases in relative income and risk of HIV infection. It may also be the case that the relationship observed is a result of the age of the data. The study was based on mortality data and the infections would have occurred some years previously. Since the data were collected in the early 1990s, the majority of the infections would have occurred in the mid-1980s. Knowledge of protective behaviours was at the time scarce and this relationship may no longer exist, as knowledge has become more widespread. This is the argument of the World Bank (1997). This is

also why further research is required. The argument put forward that, as knowledge of protective behaviour options became more available the pattern of infections will change, is very plausible. The question, however, remains as to what the relationship will change to. The more educated will have reduced risk as a result of better access to information, but does this mean that the poorest, with the least access to information, are now the most likely to be infected?

4.2.2 Country studies

Studies examining differences in infection rates between countries are more likely to measure the role of absolute income in the spread of HIV, as opposed to individual level studies, which are likely to pick up the influence of relative income. This is not to be interpreted to mean that an entire country can be considered as a single reference group, but rather that, other things being equal (especially income distribution), the higher the average income of a country, the fewer the number of people living in poverty: the fewer people in poverty, the greater the protection against infection resulting from absolute income. In other words, in a country comprising many reference groups, other things being equal, the higher the average income of the country the fewer of the countries reference groups living in poverty. The fewer of the countries reference groups in poverty the lower the average risk of all reference groups and therefore the lower the prevalence of the country. This lower prevalence is the result of the protection against infection associated with a move out of poverty – the absolute income affect.

The Over (1998) study, which was outlined in the literature review, examined the influence of a range of variables on the infection rates measured in a number of countries. One of the variables examined[9] was the GNP per capita. This variable yielded a negative and significant coefficient across all the regressions reported on in the paper. This result fits with the theory under examination, supporting the hypothesis that absolute income offers protection against HIV infections.

[9] The method of examination is discussed in chapter 1 and is not repeated here

4.2.3 Summary

The two studies discussed in this section at first glance appear to contradict one another. The individual study identifies a positive link between HIV and socio-economic status, while the second country study finds a negative relationship. Previous efforts to explain such findings have argued that the contradiction is related to the age of the data (World Bank, 1997). While it is plausible to say that the role played by income in the spread of HIV changes with time, this is a partial, not an adequate, explanation for the apparent contradiction, as it does not explain what it might change to. Although the World Bank presents data on how the pattern has been changing, it never endeavours to answer the question as to how far the change goes. Now that the epidemic has matured in Africa and information is available on protective measures, will it never be the case that individual studies identify a positive, and group studies a negative, relationship between income and HIV infection? In an effort to answer this and other questions relating to what role income plays in a mature epidemic, this thesis has attempted to build on the earlier explanations, by distinguishing between absolute and relative income. The approach taken makes it simpler to distinguish between the different influences of income on risk of infection and to examine these at different stages in the epidemic. The theory developed here is, therefore, not an attempt to contradict earlier work, but rather to develop a more in-depth understanding of the issues at hand. The theory provides a framework within which the apparent contradiction can be understood and explained.

4.3 Antenatal clinic data

The antenatal clinic data allowed for the linking of individual HIV status with a number of individual characteristics and census data. The following discussion focuses on the variables included in the final two models. A list of the other variables examined is included as an appendix. Two models were estimated, one assuming that the entire sample was the most appropriate reference group proxy. The second assumed the local community to be the appropriate reference group and identified women's relative

position in that group by comparing their education to that of other attendees at the same clinic. Table 4.1a presents the results of the first model and Table 4.1b the results of the second:

Table 4.1a Logit regression results assuming the entire sample as the reference group

Dependent variable	Coefficient (P>z)
Age 15-19	-0.459* (0.000)
Age 20-24	0.174** (0.078)
Age 25-29	0.241* (0.017)
Age 35-39	-0.344* (0.028)
Age 40-44	-0.629* (0.040)
Years of Education	0.157* (0.000)
Years of Education ^2	-0.009* (0.002)
Female status	-0.001* (0.011)
Male short distance work related mobility	1.69* (0.005)
Monthly Household Expenditure	0.00013* (0.004)
Constant	-1.631* (0.000)
Log likelihood of fitted model	-2782.32 ($P>chi^2$ 0.000)

*Significant at 5 per cent level

**Significant at 10 per cent level

Table 4.1b Logit regression results assuming the local community as the reference group

Dependent variable	Coefficient (P>z)
Age 15-19	-0.449* (0.000)
Age 20-24	0.183 ** (0.065)
Age 25-29	0.241* (0.017)
Age 35-39	-0.345* (0.028)
Age 40-44	-0.628* (0.040)
Years of Education	1.314* (0.000)
Years of Education ^2	-0.607* (0.000)
Female status	-0.001* (0.008)
Male short distance work related mobility	1.701* (0.005)
Monthly Household Expenditure	0.00012* (0.003)
Constant	-1.610* (0.000)
Log likelihood of fitted model	-2781.9 ($P>chi^2$ 0.0000)

*Significant at 5 per cent level

**Significant at 10 per cent level

A cursory comparison of the two models reveals that the different assumptions make little difference to the results. Ninety five percent confidence intervals for each variable,

for both models, were calculated and the estimated values for corresponding variables from each model, are all within the 95 per cent confidence interval of the same variable. No coefficients have different signs when the models are compared. The focus of this analysis is on the signs of estimated coefficients; as they do not differ between the models the discussions would also not differ if they were based on the results of one or the other model. The difference between the results of the two models will, therefore, only be discussed in relation to the education variables. The remaining variables will be discussed with reference to the first model. The discussion on the education variables will be left until last.

4.3.1 The age variables

The distribution of HIV between the age groups was addressed through the inclusion of five dummy variables. The women attending the clinics were grouped into six age groups. These groups were defined as five-year age bands beginning with 15-19 year-olds. For the purpose of the regression, the age group 30-35 years old was omitted and used as the base case. At the 5 per cent level, all the variables were significant with the exception of the 20-24 year old age group, which was significant at the 10 per cent level ($p=0.078$).

While necessary to control for the effects, the issue of age and risk of infection has not been a focus of this work, as it introduces additional complexities. Indeed, some of the effects of income may be captured as older women may have more experience and therefore higher earning potential. Disentangling such effects would not have been possible with the current data. It is even more difficult to interpret the age coefficients here in relation to the theory, as there are a number of biases in antenatal clinic data as representative of all women. At young ages, antenatal clinic data overestimate prevalence. Young women attending clinics are, on average, more sexually active than other women of the same age who are not attending. Since many women are still not sexually active at all at young ages, those who are pregnant are clearly more active. This bias largely disappears for older age groups as a larger portion of women become

sexually active. HIV, however, reduces fertility. Among the older age groups, therefore, other things being equal, HIV prevalence among women attending antenatal clinics underestimates the prevalence of women not attending. Although difficult to interpret, the results do exhibit the same pattern as that observed in antenatal clinic data on HIV throughout Africa. HIV prevalence increases between 15 and 29 years of age and then decreases (Stover, 1999).

4.3.2 Female status variable

The female status variable was calculated as the average income of women in the magisterial district in which the clinic was located, multiplied by the ratio of average female income to average male income. It is the combination of female poverty and inequality between the sexes that is argued to substantially increase risk. In chapter 2 it was argued that, other things being equal, the relative income effect would be strongest at low levels of female income, when women are vulnerable; the greater the inequality in income the greater the ability of men to exploit that vulnerability. The combination of female poverty and sex-based inequality has been argued, in chapter two, to increase the risk of the entire group; firstly, by increasing the risk of vulnerable women, secondly, by increasing the risk of the relatively wealthy men who have power over them and finally of the other partners of both the vulnerable women and the relatively wealthy men. This variable captures both the absolute level of women's status and their status relative to men in the magisterial district. The variable is lowest when women have low average income and the differential in income between men and women is large. The variable is highest when the average income of women is high and there is little difference between male and female income. According to the theory, the lower the average economic status of women in the community and the greater the sex-based inequality, the greater the risk of infection for everyone. The expected sign was, therefore, negative. That is to say an increase in women's average income, or in the ratio of female to male income, or a combination of the two, decreases the risk of infection for the entire group.

The result observed was, as expected, negative and significant at the 5 per cent level with a p value of 0.011. For a 25-30 year old woman, all other variables at their mean, a shift from the lowest female status observed to the highest decreased risk by 0.1 from 0.39-0.29.

4.3.3 Male short distance work-related mobility

A number of measures of mobility were examined during the analysis, including variables indicating the level of long distance work-related movement. Long or short distance work-related mobility was distinguished as follows: short distance mobility refers to work-related mobility within the province, and long distance between provinces or countries. Variables that captured both male and female mobility were also examined. Male short distance mobility was finally settled on because of the strength of the relationship with risk of infection. The other variables examined, which are described in appendix A, were not significant, either as the only group variable or as part of the final model. The variables examined were constructed from the census data based on the magisterial district of households with employed members and the magisterial district of those employed household members' place of work. Household members whose workplace was located in a magisterial district outside of the province in which the household that claimed them as members is located, were considered as part of the long distance mobile population. Members whose workplace was in the same province but a different magisterial district from their household were considered as part of the short distance mobile population of that area. The variables were calculated as the size of the mobile population as a percentage of the appropriate magisterial population, either all working age adults or working age adults of the same age, and sex.

At the onset of an HIV epidemic, long distance mobility plays an obvious role in the introduction of HIV into communities. As the epidemic progresses, the role of long distance mobile populations is likely to become less important. Those travelling long distances for work typically can be expected to return home infrequently, when compared to those travelling shorter distances. If the number of times an individual returns home is

a function of the cost of returning home, which in turn is a function of the distance, the further away a person travels for work, other things being equal, the less frequently they return home. At an advanced stage of an HIV epidemic, both short distance and long distance mobile populations face the possibility of coming into contact with HIV infected sexual partners while away from home. The short distance populations, however, are returning home more frequently (due to the nature of the variable's construction it includes some individuals who return home daily) and are more likely to infect people in their own community. The data under examination are from two high prevalence provinces and, while measures of long distance mobility of male or female populations were not significant, measures of short distance mobility were.

The short distance male, work related, mobile population variable was significant at the 5 per cent level, with a p value of 0.005. The variable was, as expected, positively related to risk of infection for women in the community. For a 25-30 year old woman, with all other variables held at their mean, an increase in male short distance migration from its minimum value to its maximum increased risk by 0.09 from 0.35-0.44.

Due to the nature of the variable's construction, some of the influence measured may be related to urbanisation. The variable was constructed from the Census as the proportion of male household members, working outside the magisterial district in which the household was located but within the same province. A number of cities in the two provinces contain more than one magisterial district and many men may travel back and forth each day; therefore, in urban areas this variable yields a high value. The principle is, however, the same. Urbanisation and mobility increase the size of sexual networks thus increasing risk, which is what the variable was intended to identify. Interestingly, urbanisation measures, measured as the proportion of households in the magisterial districts which were defined as being in an urban area, on their own were not significant, even when the short distance migration variable was omitted from the model.

4.3.4 Monthly household expenditure

This variable is important, but difficult to interpret in relation to the theory developed in this thesis. It is the average monthly household expenditure of all households in the area, not the subset of individuals defined in the previous chapter and used in the calculation of other census variable averages, that is used in this analysis.

This variable was not intended as a proxy for absolute income because, as discussed in the previous chapter, there are unlikely to be large variations in absolute income within the sample. All the women tested were attending state antenatal clinics; such a similar level of health care suggests that absolute income levels are fairly similar. Not only is the level of absolute income of the women likely to be fairly similar, it will also be lower than the average absolute income of all women in the same magisterial district. This will be the case because sampling from only state clinics means that wealthier women who would attend private facilities are excluded. Furthermore, the average status of women in the community has already been controlled for. This variable, therefore, is intended to capture the influence of regional economic characteristics on the poor (women attending state clinics) living in that region. Areas with high household expenditures are likely to be wealthier areas, attracting more labour migrants, traders and shoppers. It was mentioned previously that the level of outward mobility is an important determinant of risk. This variable, however, seeks to capture the influence of inward mobility on risk. The variable was, therefore, expected to yield a positive sign as wealthier areas are likely to have higher levels of inward mobility and as a result higher levels of risk. The results expected here must be clearly distinguished from those associated with absolute income. The expenditure variable is used in this context to measure the influence of different regional characteristics on groups with similar absolute income. That is to examine the difference in risk for poor women living in poor areas, compared to poor women living in wealthy areas.

The coefficient was positive and significant at the 5 per cent level with a p value of 0.004. This suggests that, as the overall expenditure of an area increases, the risk of its poor (women attending state clinics) increases. For a 20-30 year old woman, all other variables at their mean, a shift in monthly household expenditure from its minimum of R800 to its maximum of R5000 increased risk from 0.31-0.43.

4.3.5 Education

For the purposes of this analysis, the results of the education variables are the most important. They are intended to proxy for relative income. Education was, therefore, expected to have a positive influence on the risk of infection. It was, however, unclear whether the measure should be relative to the entire sample or to other women attending the same clinic. Models were estimated under both assumptions and the results appeared to be insensitive to the choice of model. In both cases, the direct measure yielded a positive coefficient and the squared term a negative. Both education variables in both models were significant. Given that the assumption made little difference to the results and the other variables have been discussed in relation to the first model, the education results will also be interpreted in the same way.

The results were as expected. As years of education increased so did risk. This increase, however, occurred at a decreasing rate. Although it has been repeatedly stated that there is not sufficient variation in income to look for protective influences of absolute income, there are two reasons why it was expected that the education squared variable should have led to a negative coefficient. Firstly, assuming diminishing marginal utility to additional partners for men, the relationship between risk and income for men would be expected to be non-linear. As the relationship between women's income and their risk is predicted, by the theory developed in chapter two, to follow a similar pattern to that of their male partners, the relationship here is also expected to be non-linear. Secondly, the small number of better educated women attending state antenatal clinics may be at the bottom end of a higher reference group. Although they have not gained access to better health care, they may have benefited from the higher level of education. Interestingly, the

point at which risk starts to decrease is similar to the point where Case and Yogo (1999) identified increasing rates of return to education.

When the education variables are examined collectively the following relationship is observed:

Figure 4.3 Probability of infection and years of education[10]

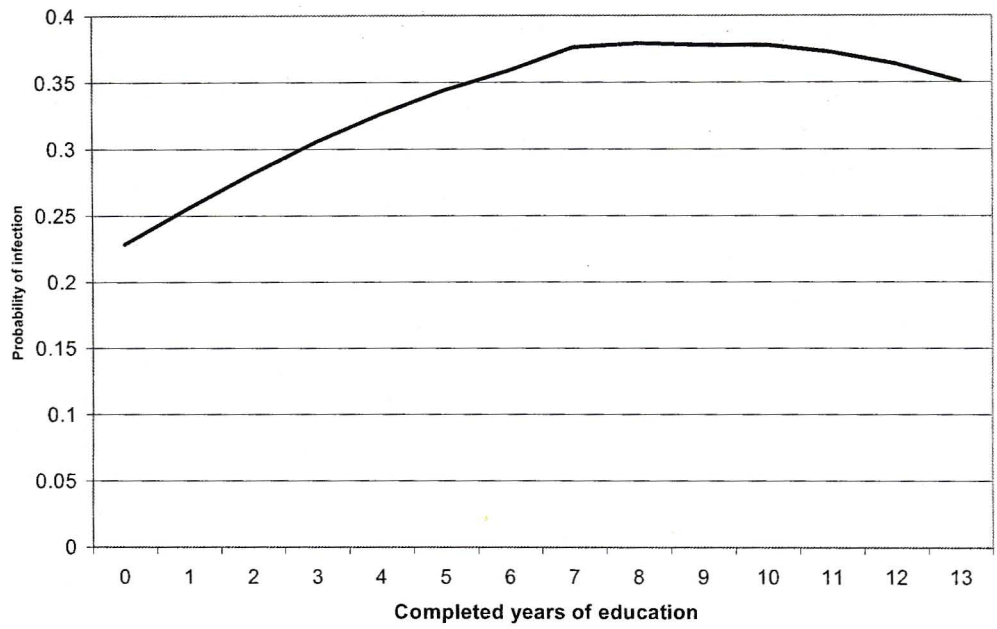

The probability of being infected increases from 0 years of completed schooling until 8 years where is plateaus. After 8 years of schooling the probability begins to decline. The sample size for 12 and 13 years was, however, very small.

A movement across the entire range changes risk from 0.23 at 0 years of completed schooling to 0.35 at 13 years. This, however, masks the non-linear relationship and the peak in risk at 8 years of completed schooling.

[10] For a 25-30 year old women, all other non-age variables held at the mean

The most important result is that the risk of HIV infection increases, to a maximum point, as educational level increases. As far as education is a correlate of income, this result can be interpreted as evidence that, in an advanced epidemic, there are still times when increases in income increase risk. This is clearly in line with the theory developed in this thesis which argues that, within reference groups, increases in relative income increase an individual's probability of infection.

4.3.6 The model

Overall, the regression results were significant at the 5 per cent level with a p value of 0.0000. The log likelihood of –2782.3224 does not represent a very close fit. Due to the nature of the data and the subject, substantial individual variation was expected and, while it was hoped that the regression would be significant, a close fit was not expected.

It is impossible to fully specify a model of this nature. There will be many other factors involved that have not been captured here. As a result there will obviously be omitted variables and the associated bias. This has important implications; firstly, the coefficients on some variables included in the model will be biased if they are correlated with a variable excluded from the model. An example of this is the over estimation of the explanatory power of education. This implication is, however, desirable. Education, in this model, is included because it is a positive correlate of income. Indeed the results are interpreted on the assumption that such a bias exists. Other coefficients may also be biased. Where this is feared to have happened it has been discussed in the text: for example, the inclusion of some of the effects of urbanisation in the coefficient of the male mobility variables.

Unfortunately, however, the nature of the logit is such that even if the omitted variables are not correlated with included variables, estimations are inconsistent. This is unavoidable given the impossibility of specifying a model of this nature completely.

Given this inconsistency and the problems associated with addressing heterosecasticity in the logit framework, both models were also estimated using ordinary least squares with robust standard errors. In both the re-estimated models the variables included in the logit model were significant at the 5 per cent level with the exception of the dummy for aged 20-24 years, and the coefficients yielded the same sign as they did when estimation was made using the logit.

4.3.7 Summary

Two models were estimated, based on different assumptions about the appropriate proxy for individuals' reference groups. Both models were significant and the results differed only slightly. As there was little difference between them, the discussion of results concentrated only on one model, the one that assumed the entire sample to be the appropriate reference group.

The age group dummy variables are difficult to interpret in relation to the theory, although they did exhibit the expected pattern. The results of the other variables are summarised in the following table.

Table 4.2 The probability variation in risk of infection for a move from the minimum to maximum value of each variable*

Variable	Min value probability	Max value probability
Education	0.23	0.35
Female status	0.39	0.29
Male short distance migrants	0.35	0.44
Monthly Household Expenditure	0.31	0.43

*For a 25-30 year old women, all other non-age variables held at the mean

The table displays the change in risk, for a 25-30 year old woman, for a change in the minimum to the maximum of each variable, while the other variables are held at their

mean. The table, however, masks the non-linear relationship between education and risk, which peaks at 8 years of completed schooling and then declines.

Collectively, the results imply that communities in which women generally have low status, and male mobility is at a high level, risk of infection is at its greatest. Interestingly, and in line with the theory, it is the relatively better off women within these communities who appear to be at the greatest risk, so far as education is a correlate of socio-economic status. The results further suggest that the lower the average status of women and the higher the sex-based inequality in favour of men, the greater the risk of the entire group. But the higher the status of an individual woman within that group, the higher her risk, again in so far as education is related to status. This result is in line with the predictions of the theory developed in chapter two. Lower female status increases the power of men in the group; this power is further increased the higher the income differential between men and women in favour of men. This in turn increases the risk of poor women, and of relatively wealthier men and their partners. This increases the risk of all women but, as discussed in the theory chapter, the increase in risk is greatest for the partners of the wealthier men, who are themselves likely to be of higher socio-economic status.

The above analysis provides strong evidence of the positive relationship between relative income and individual risk of infection. It does little, however, to address questions of absolute income. The following section seeks to deal with these issues in more detail.

4.4 Analysis of Company X data

This section discusses the results of the regression analysis conducted on data collected from a large South African company, the data consist of observations drawn from across the country and from all job bands within the company. The important variables are those associated with income; these are based on the average income of individuals with the same characteristics. The method used to generate the income variables is outlined in the previous chapter and the appendix of this thesis. The income variables are designed to

address issues relating to absolute income. Expectations based on the theory developed in this thesis are that, across a large range of income, the risk of infection will decrease as income increases, but at a decreasing rate. As an income and an income squared variable were included in the estimation, it was expected that the simple income variable would generate a coefficient with a negative sign and the squared term a positive sign. The simple term measures the direction of the relationship and the squared term the rate at which this relationship weakens. The following table summarises the results of the regression.

Table 4.3 The results of the logistic regression on company data

Dependent variable	Coefficient (P>z)
Aged under 25	-0.399 (0.371)
Age 25-30	0.006 (0.981)
Age 35-40	-0.132 (0.564)
Age 40-45	-0.457** (0.054)
Aged over 45	-0.751* (0.001)
Eastern Cape	-1.159* (0.014)
Free State	-0.407 (0.125)
Gauteng	-0.680* (0.012)
KZN	0.629* (0.001)
Northern Cape	-0.035* (0.930)
North West	0.278 (0.267)

Northern Province	-0.348** (0.063)
Western Cape	-1.38* (0.041)
Asian	-0.686 (0.163)
Coloured	-1.247* (0.019)
White	-1.393* (0.000)
Single sex hostel	0.349** (0.067)
Construction camp	0.294 (0.102)
Other accommodation	0.267 (0.159)
Other accommodation	-0.589 (0.219)
STI	0.489 (0.020)*
Income	-0.0000357 (0.000)*
Income squared	1.25E-10 (0.000)*
Cons	-0.665 (0.878)
Log likelihood of fitted	Log likelihood = -966.657 (P>chi^2 0.000)

*Significant at 5 per cent level

**Significant at 10 per cent level

The analysis included variables relating to employees' age, ethnic group, the location of their work place, their accommodation, STI history, and income. Although not all are directly relevant to the role of income, the results of each variable are discussed as they have a role in determining an individual's risk of infection.

4.4.1 Age variables

As in the previous analysis, the age variables are difficult to interpret, though for different reasons. These data are based on a male sample, so they do not have the bias associated with fertility. Nor do the data suffer from bias introduced in antenatal clinic data for the young by sampling only the sexually active. Indeed the pattern is as would have been expected for a male distribution, although the differences here are largely insignificant. Typically, the peak of male infection is later than that for females. Although not significant, this pattern can be seen here if the data are compared to the results in the previous section. It would be reasonable to expect that age, power and income are linked and that the age distribution would be an important part of this examination. It would, however, also be reasonable that the links between these variables would be different between reference groups. As this sample draws from many reference groups, interpretation of the results along such lines is not possible. What is interesting, however, is that there is little significant difference between the age groups. The omitted age group is the men aged 30-35 years. With the exception of the over 45 year olds, the dummies for all other age groups were insignificantly different from the omitted case at the 5 per cent level, although at the 6 per cent ($p=0.54$) the 40-45 year old age group was different. This lack of significant differences may be linked to the inclusion of income in the model. Much of the risk associated with older men is derived from their power in their reference group, which is also likely to be linked to their income. In other words, much of the age distribution may be associated with age as a determinant of income patterns, so once income is included in the model the differentials are weakened. Despite their probably higher income, older men still display a significantly lower prevalence. This is partly a result of a higher average age of their infection, that is, the length of time since they were infected. By that age many of the members of such an age cohort would have already died as a result of HIV progressing to AIDS.

4.4.2 Provincial site location

The company at which these data were collected has operations across the country. This set of dummy variables indicates the provincial location of the site where each individual works. In South Africa there is substantial provincial variation in prevalence. There are explanations as to why this is the case (see Dorrington and Johnson, 2001 for example). Although a closely related topic, this is beyond the scope of this thesis. It is, however, necessary to include provincial dummy variables because a major determinant of an individual's risk of infection is the prevalence of the communities in which they live.

Mpumalanga was selected as the omitted case as more observations were drawn from there than any other province. KwaZulu-Natal and the Western Cape were significantly different from the base case - both with the expected sign, given the current antenatal clinic rates in the provinces. The Northern Cape and the Northern Province coefficients yielded the expected sign but were not significant. The North West did not display the expected sign but was also not significantly different from zero.

4.4.3 Ethnic group

The issue of ethnic group and HIV is a sensitive one. Although related to the topic of this work, it is not directly relevant to the discussion. It is, however, worth making a few comments given that the variables are included in the model. Black was used as the base case, with Asian, Coloured and White included as dummy variables. The coefficients on all three dummy variables were negative, but only the White and Coloured variables were significant. The Asian variable was probably not significant because there were very few observations.

There are four primary reasons for ethnic group variations in prevalence in South Africa. The first is income, where the differentials are vast and well known. Efforts have been made here to control for the income component of the variations. The second and closely related issue is that of education. On average, whites have received better education. This offers them a great deal of protection, the same as that discussed in relation to absolute income. Thirdly, as a result of the Apartheid system, black people were often forced to live very far from their place of work. Migration, as has been mentioned in earlier sections and chapters, increases risk of infection. Finally, HIV is concentrated in Africa and was introduced into the general population of South Africa via migrants entering the country and returning from work outside. As has been discussed, HIV is concentrated within the poorer segments of society. The communities into which migrants from these high risk groups moved tended to be black. Similarly, the communities which black South African migrants moved to were the high-risk communities in those countries.

4.4.4 Accommodation

Data were collected on the type of accommodation employees lived in: their own house, a rented house, a single sex hostel, a construction camp, or other. These variables are highly correlated with income, in that the low income earners are more likely to live in the single sex hostels and the construction camps. Single sex hostels and construction sites are typically associated with high risk sexual behaviour (Campbell, 1997) and living in them was expected to be associated with greater risk of infection. The results were as expected: employees in single sex hostels were most at risk, followed by those in construction camps, then rented households and the lowest risk being associated with those living in their own house or some other accommodation. None of the results, however, were significant at the 5 per cent level. This is most likely the result of the high degree of correlation with income. Living in a construction camp or a single sex hostel was correlated with income, with a correlation coefficient equal to 0.43.

4.4.5 STIs

STIs have been discussed many times as a co-factor in the spread of HIV. As expected, symptoms of an STI were positively related to the probability of infection, and the result was significant at the 5 per cent level. This probably over-estimates the risk associated with STIs, as not only are people with STIs at greater risk of infection with HIV, but people who are HIV positive are also at greater risk of infection with an STI.

4.4.6 Income

The inclusion of STIs and accommodation variables complicate interpretation of the results. They were included, however, because they allowed the investigation of different aspects of the protective role of absolute income. This protective role discussed earlier is partly attributable to better access to health care, which would imply lower STIs and better housing. Part of the impact of absolute income is therefore captured in these variables. Indeed income, STIs, and risk-taking sexual behaviour are highly correlated. The data on nights away from home and non-regular partners were not included because they were too highly negatively correlated with income and accommodation to attribute a separate influence. Collectively, STIs, accommodation and income accounted for over 50 per cent of the variation in risk taking behaviour. The results, therefore, of the accommodation and STI variables must be considered in conjunction with the income variables to determine the impact of absolute income.

Both the absolute income variables - income and income squared - were significant at the 5 per cent level, both with p values of 0.000. The income coefficient exhibited a negative sign and income squared a positive sign. Income, therefore, was found to reduce the risk of infection, but at a decreasing rate and eventually income squared becomes the dominant influence and an increase in income increases risk. The results are displayed graphically below as the predicted risk of infection allowing only income to vary, and by association income squared, all other variables held at their mean.

Figure 4.4 Income and predicted risk of infection

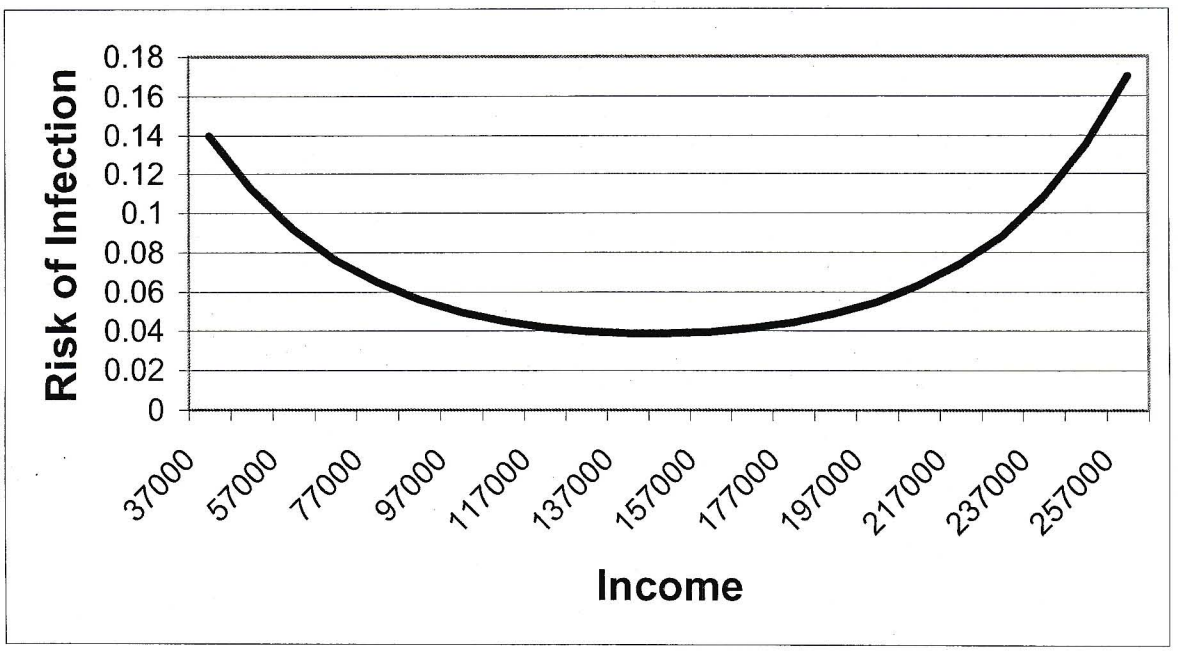

The relationship outlined above does not suggest that employees with very high income are the most at risk, as the other factors associated with income have not been included. Furthermore, it must be noted that the sample size in the higher income groups is small. The estimation process would have attached more weight to the majority of observations at lower levels of income. At these levels, the influence of income squared would be to reduce the rate at which income was providing protection. This suggests that there is a problem associated with an omitted variable. Should income cubed have been included in the model it would have been expected to yield a negative co-efficient, which would have had the effect of flattening the right hand side of this graph. Instead of including this variable, and in an effort to provide more meaningful results, the protective influences associated with income were directly included, namely differing rates of STIs and better housing. Once these variables are included, the relationship between income and risk is as expected.

The risk pattern changes dramatically when other variables are changed. A 25-30 year old black male, working in KwaZulu-Natal, living in a single sex hostel, with an STI and

earning a salary at the bottom of the range has a predicted risk of infection of 0.42. An older white male, on the other hand, working in the Western Cape, living in his own home, having no STI and earning the top salary has a predicted risk of infection of 0.01. The difference is not purely associated with race; if the top earning male were African his predicted risk would be 0.04. STIs, accommodation and location are clearly very important. For example, an African male working in Mpumalanga, with no STI, earning the lowest wage, but living in his own home has a risk of infection of 0.14. If he is infected with an STI, this increases to 0.21; if he moves to KZN his risk will increase to 0.34, and if he stays in a single sex hostel when he arrives in KZN the risk increases further to over 0.42.

Absolute income, once STI prevalence and accommodation have been considered, clearly offers protection against infection, although at a diminishing rate. This result is displayed graphically below. The points represent the average predicted probability of infection for individuals with the same value for the income variable.

Figure 4.5 The effect of income on probability of infection where STI prevalence and accommodation vary

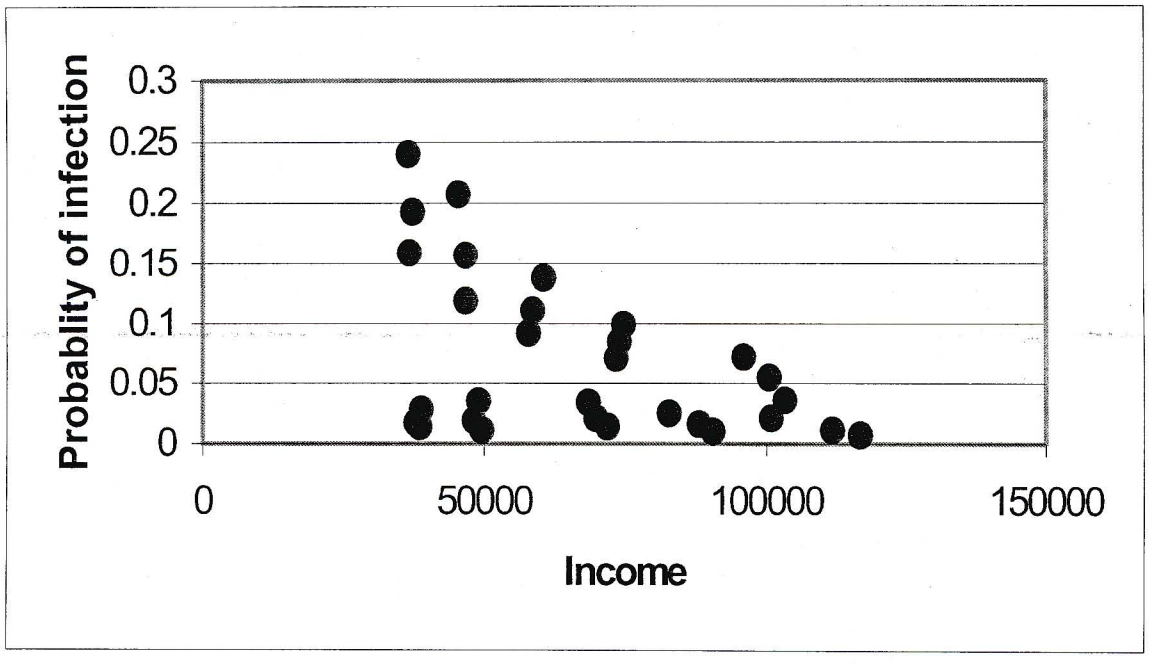

The above figure shows that once variation in accommodation and the associated behaviour, and STI prevalence are included, absolute income offers protection against HIV infection. This provides considerable support to the theory on the role of absolute income in the spread of HIV.

4.4.7 The model

Overall, the regression results were significant at the 5 per cent level with a p value of 0.0000. The final log likelihood was -966.657; although not a very close fit, the model was significant. As with the previous analysis, the nature of the data is such that, while it was hoped that the regression would be significant, a close fit was not expected.

The available data did not allow for the model to be fully specified. As a result of this limitation, there will be omitted variables. The implications of this are two fold; firstly, as with most regression models with omitted variables that are correlated with included variables, bias is introduced into the estimation by causing the included variables to be correlated with the error term. Secondly, in logit regressions there are negative implications, even if the omitted variables are not correlated with any of the included variables. In this situation, although the omissions do not bias the results, the estimation is inconsistent.

As these problems are unavoidable, given the nature of the topic and the limitations of the available data, the model was re-estimated using ordinary least squares with robust standard errors. While not solving the problems of omitted variables, the exercise did show that the results were robust to model choice. The results of the re-estimation yielded the same signs and similar levels of significance.

4.4.8 Summary

The analysis of company X data yielded the expected results. The age variables displayed the same pattern as is typically associated with male infections. The direction of the differences between the provinces, where significant, were similar to those recorded in the national antenatal clinic surveys. The differences between ethnic groups were as expected.

Income appeared to offer protection against HIV infection, but at a diminishing rate. Once other variables associated with income are also allowed to vary, the protective role of absolute income became far clearer.

The results in this section also bear some resemblance to the discussion in the theory chapter about the likely relationship between income and HIV prevalence among a cross section of the population. This discussion concluded that HIV prevalence would be a cubic function of income. In that case, income would have a positive coefficient, income squared negative, and income cubed positive. In the company X, cross section income has a negative coefficient and income squared a positive one. The results here are not a contradiction of the cross sectional discussion. This sample is truncated in so far as only employed persons are included. It is this omitted portion of the population where the positive relationship with income would most likely be observed, as it includes those living in abject poverty through to the unskilled labourers. If this part of the population had been included, the first expected effect of income (moving from low to high) would have been to increase risk, then to decrease risk but at a decreasing rate. Income would therefore be expected to have a positive sign, income squared a negative sign and income cubed again positive. Starting at a higher income level, as is the case with company data, the first expected effect of income would be to decrease risk, but at a decreasing rate. Income would therefore be expected to show a negative relationship to risk and income squared a positive.

4.5 Analysis of further company data

The data from company X allowed for the investigation of a number of different issues associated with income and the risk of HIV infection. For example, it allowed for the separation of the role of income from that of ethnic group. Data available from other companies are not as detailed and do not lend themselves to such analysis. They do, however, provide a view of the unadjusted relationship between income and HIV prevalence.

The following are data from two other companies in the region that have conducted surveys among their employees. This figure displays the results, by job band/level for the two companies and provides the data from company X as a comparison.

Figure 4.6 HIV prevalence by job band

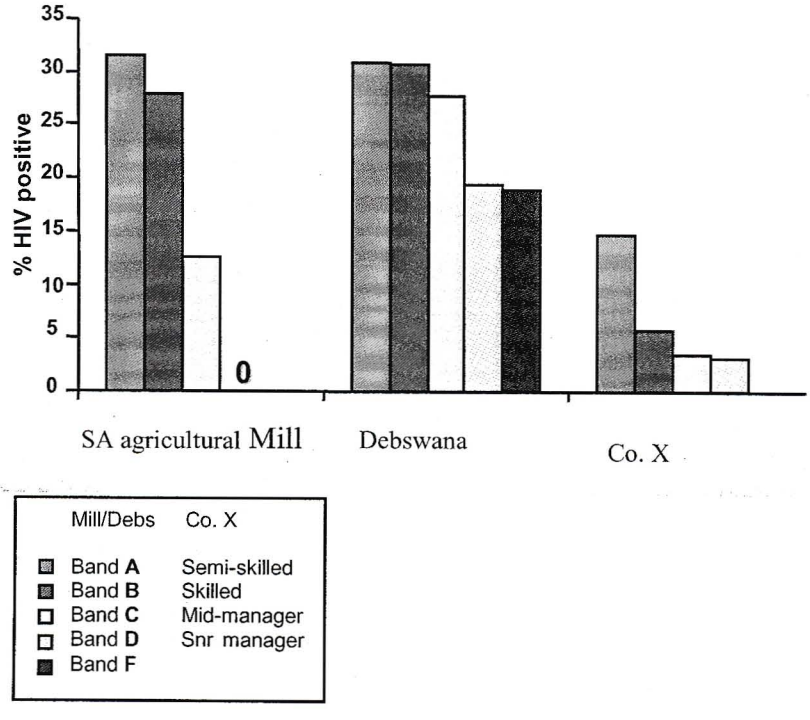

In both situations the prevalence falls as the level of employment increases. As higher job bands are associated with higher income, the graphs indicate falling prevalence with increasing income. The observed decrease in the prevalence in the South African agricultural company data, however, cannot be interpreted solely as the protective influence of absolute income, as the composition of the job bands with regard to ethnic group has not been considered. For both companies, the graphs also represent different age structures across job bands and associated differences in prevalence. These issues aside, the data provides further evidence that HIV prevalence decreases as average income increases.

4.6 Summary

This chapter was divided into four section: the first re-examined previous studies, the second analysed antenatal clinic prevalence data, the third data from company X and the final section looked briefly at data from other companies. The results of these analyses are summarised as follows:

The first section revisited the findings of two types of studies: an example of an individual risk study and one of a study conducted at a country level. On the one hand, individual level analyses have typically found that the risk of HIV infection increases as an individual's socio-economic status increases. On the other hand, country level studies, have found that as the income of a country increases so the prevalence rate falls. Efforts to explain these apparently contradictory findings have argued that the individual level studies were measuring the relationship between the risk of infection and income, using data on infections that had occurred at the early stages of the epidemic. At these early stages of the epidemic, it has been argued, knowledge on how to modify behaviour in order to protect oneself was scarce. As this knowledge became available, the wealthy would be the first to change their behaviour, as they have better access to information. The theory presented in this thesis expands on this argument in an attempt to develop a more comprehensive understanding of how such a contradiction could have occurred and to understand how these findings may have changed as knowledge has become more

widely available. While previous work focused on how changing information would change the role of income, this thesis goes further to explore what this changed role might be. The findings of the two studies revisited in this chapter are in line with expectations based on the theory developed in this thesis. Individual studies are argued to find primarily the effects of relative income, while country studies measure the role of absolute income. Within the framework of such a theory, the results of the studies were not in contradiction when they were conducted nor would they be in contradiction if they were conducted now and yielded similar results. While the availability of information will have almost certainly changed the distribution of infections within a population, the theory developed here suggests that there will still be times where a positive relationship between income (relative at least) and risk of infection will be found.

The studies outlined in the literature review and the examples re-examined in this chapter were not designed to test the hypotheses posed in chapter 2. For this purpose additional analyses where conducted using antenatal clinic and company data. The results of the antenatal clinic data analysis suggest that even at an advanced stage of an epidemic there is a relative income effect. That is to say that there are times when income is positively related to risk. This provides support for H_2, that relative income increases an individual's risk of infection, at least in regard to women, although the data did not allow for the examination of H_2 in relation to men. The argument that relative income increases the risk of women was a difficult one to make compared to the same argument for men and is based largely on the assumed link between women's risk and the risk of their partner. Therefore, while the findings of the analysis do not provide direct support for H_2 in relation to men it does so indirectly.

The theory of the dual role of income in the spread of HIV does not suggest that income will always be positively related to risk of infection. Indeed, this is argued only to be the case within reference groups. The theory predicts that, when individuals are taken from different reference groups with a wide range of absolute income levels, the trend will be a decline in risk as income levels increase. Within such a sample there may well be individuals from the same reference group; for this sub sample, risk would be expected to

increase for those with higher income levels, but across the range of income the risk of infection would be expected to fall. This expectation, the basis of H1, that as absolute income increases so risk decreases, was examined with the use of company data. Indeed the results of the analysis were in line with H1. As income increased, risk was found to decrease, albeit at a decreasing rate. A similar trend was observed in HIV prevalence levels at different job bands from two other companies, although no controls could be used to isolate how much of the trend was a result of increasing income.

Overall, this chapter has presented results which are in line with expectations based on the theory of the dual role of income in the spread of HIV. Two data sets were examined, both collected in the same country at a similar time. One suggests that as income (education) increases, so does risk, while the other suggests that as income increases risk decreases. These results have been generated using data from a mature epidemic. The theory developed in this thesis provides a possible explanation of these results, which argues that far from being contradictory, the results are as expected. The data sets examined where, as discussed in each section, limited. The analyses are based on a number of assumptions and the findings here are only in line with the theory in so far as those assumptions are in line with reality. While it is argued that all these assumptions are reasonable, as with any analysis of this type the results always need to be treated with some caution.

Chapter 5

The Extension of the Theory of the Dual Role of Income in the Spread of HIV

5.1 Introduction

The theory and the analysis thus far have examined the role of income in the spread of HIV in terms of a static analysis. The emphasis has been on what the differences in risk are for people with different incomes. This chapter develops the theory to include a dynamic component and discusses what the likely implications are for changes in individual income. This chapter also deals in more detail than before, with the role of different income distributions.

While the results outlined in the previous chapter support some of the hypotheses developed in this chapter, the data are largely insufficient to test the aspects of the theory developed here. Although the testing of the theory in this chapter is not possible at this stage, it is useful and interesting to apply the concepts developed more broadly.

5.2 Individual changes in income

A change in the level of income of an individual changes both the average absolute income of the reference group and also the individual's position within the group. The resultant change in risk will be determined by a number of factors. It may also be the case that the change in risk for an individual with income X resulting from an increase in income to income Y might be different from the difference in risk between an individual with income X and an individual with income Y.

To examine the implications of a change in income, it is useful to follow through the likely impact of an increase in income for a male and a female.

Consider first a male with a given income. His risk is determined by the absolute income of his reference group and his relative income within that group. If his income increases, this increases the absolute income of his reference group as well as his relative income within the group. Unless his increase in income is very large, or his reference group is very small, the resultant increase in absolute income of the group will be small and unlikely to make any difference to the group's risk. His increase in relative income increases his power within his reference group and his attractiveness, these two factors work together, and depending on his sexual preferences, will increase his risk of infection if he demands more sexual partners, or leave it stable if he does not. On average, therefore, an increase in a male's income will increase his risk of infection. This is the same relationship as that discussed in the cross section; a decrease is also likely to be similar to the cross sectional relationship.

The influence on risk of a female change in income is more difficult to determine. According to the principle of positive assortative matching, the higher the socio-economic status of a woman the higher the likely status of her partner. Therefore, if the income of a female increases, she will likely partner with a male of higher socio-economic status. The risk that her partner is, or will become, HIV positive is therefore greater, thereby increasing her risk of infection. This again is the same relationship as in the cross section. It fails, however, to deal with women who are already in a partnership. An increase in income for a woman already partnered increases her power within the relationship. This increase in power can reduce her risk of infection. For example, a woman with greater power in a relationship will be more able to influence condom use and possibly the number of other partners her partner has; both actions reduce her risk. The influence on the risk of infection of an increase in income is different for women with partners compared to those without. Similarly, a decrease in the income of a woman who is already partnered decreases her power within the relationship, thereby, on average, increasing her risk of infection.

Changes in an individual's level of income affect their relative income and therefore their risk, but the influence on absolute income is likely to be too small to have any meaningful effect on risk. In contrast, if the income of every member of the group increases, then absolute income could be affected to such an extent that the risk of the group changes. There are, therefore, changes that affect the group's average absolute income to the extent that they affect the risk of the entire group but do not affect the relative income distribution, while others only change an individual's relative income. Changes in the income of sub-groups of individuals within the reference group or changes in the distribution of income within the group lie in between these two types of changes, affecting relative income and, where large enough, absolute income. Both these changes in distribution and of sub group income are discussed in the following sections.

5.3 Income distribution

Income distribution can be discussed in two contexts: 'within group' income distribution and 'between group' income distribution. 'Within group' income distribution is the driving force behind relative income. If income were distributed equally within each group, every individual would have an income equal to the group average and there would be no relatively deprived or wealthy individuals. Within groups, two types of income inequality are important in determining risk: between sex inequality and within sex inequality. The greater the income inequality between men and women the greater the power men have over women. If, as it is assumed, on average men prefer more partners to fewer, the greater the power men have the more sexual partners they will have, placing themselves and their partners at greater risk of infection. An increase in the income inequality between men and women in favour of men will, therefore, increase the risk of the entire group. A distribution of income towards women, while maintaining the same average income for the group, would reduce the risk of the entire group by reducing the power of men within that group.

In addition to income inequalities between the sexes, inequalities within sexes are also important, particularly for males. Core groups are important in the spread of HIV.

Income inequality among men creates a core group of relatively wealthy men. This core group is arguably responsible for spreading HIV among the poorer women of the reference group and indirectly to the poorer women's male partners. If the power relationship between men and women is in part a function of the absolute level of, and the relative difference in, income, then, as male inequality increases, the power of the men at the upper end of the distribution will likely be increased, increasing the likelihood of risky behaviour. At the same time as inequality increases the power of the men at the upper end of the distribution, it decreases the power of those at the other end. This is, however, unlikely to negate the entire increase in risk resulting from the increase in the power of men at the other end of the distribution. The increase in power, it is argued, increases the risk of the men at the top end of the distribution and by a similar margin the risk of their partners but it also, albeit to a lesser extent, increases the risk of women at the lower end of the distribution. These women are likely to partner the men at the lower end. The increase in men's power, therefore, could well increase the risk of all men and women in the reference group.

Income inequality among females is also important. An increase in income inequality will likely increase the power of men over women who lose in the redistribution. At the same time, the risk of women who gain may be increased, as they are likely to partner with higher risk men, if they have not already partnered.

'Between group' income inequality is also likely to play a role in determining risk. Arguably, other things being equal, a population with a given average income will have higher rates of infection the higher the rate of income inequality. Increases in absolute income should offer protection against HIV but, except at very low levels of absolute income, these increases are likely to lead to progressively smaller decreases in risk. A redistribution of income from the wealthy should, therefore, result in a smaller increase in risk of infection than the decrease in risk resulting from the increase in income of the poorer members of the population.

In the context of this theory, income inequality is argued to affect the spread of HIV in a number of different ways. Inequality between the sexes is likely to influence power relations and therefore risk. Inequality within the sexes affects the distribution of power and therefore risk. Inequality between groups affects risk as a result of the non-linear relationship between absolute income and risk of HIV infection.

The discussion of the role played by the different types of inequality can be combined with the analysis of income changes and their influence on individual risk of infection, in order to examine the implications for reference groups of changes in the income of sub groups – that is, groups of individuals within reference groups as opposed to sub reference groups.

5.4 Sub-group changes in income

Some sub-groups can change income levels and the net impact on group risk is unclear, while for other groups all factors work in the same direction and the net influence on risk is more obvious. For example, if a group of males gains some additional income, male inequality may increase and male female inequality will increase. If both these increases in inequality occur they will increase the risk of the entire group. If, however, the group and the increase in income are large enough there may be some protective influence associated with an increase in the references group's absolute income. The two factors work in opposite directions, thereby making the net outcome on risk of infection indeterminate. If the increase in income increases male:female inequality but decreases male inequality then the inequalities will also work in opposite directions, further clouding the outcome.

If the sub-group were all women, the situation would also be complicated. The increase in income could increase the risk of the women who gain or, as discussed previously, it might reduce their risk if they are already in a more permanent partnership. Even if it did increase their risk, the increase is relative to other women by improving their socio-economic status and thereby increasing the likely socio-economic status of their partner.

But just as this increases their risk, it reduces the risk of women who may now partner with a male with lower socio-economic status because they were displaced by the women whose income increased. The increase in relative income is therefore likely to result in no change in average risk or a reduction in average risk. At the same time, the increase in female income will reduce the inequality between men and women and in so doing reduce the power of men in the reference group and reduce the average risk of the group. If the group and the increase in income are large, the increase in income may be large enough to improve absolute reference group income enough to benefit from its protective influence.

If the sub-group were mixed, the net effect of an increase in income would be very difficult to determine. The more income that goes to the relatively deprived, regardless of their sex, the more likely the net outcome is to reduce risk. Similarly, if income increases, the greater the proportion of that increase that goes to women, the better the net outcome for the entire group's risk will be.

5.5 Summary

This chapter has attempted to extend the Theory of the Dual Role of Income in the Spread of HIV developed in chapter 2. Prior to this chapter, the focus had been on explaining the reasons for cross sectional variations in HIV prevalence that are in part a result of different income levels. This chapter has introduced a dynamic component, discussing how changes in income for groups and individuals might change their risk of infection.

What is important to note from the chapter is that income distribution and individual and group changes in income have implications for the risk of the entire reference group. Furthermore, increases in income that affect absolute income may not always have a positive outcome if the distributional impacts worsen the situation more than the increase in absolute income betters it.

Some of the theory developed in this chapter does relate to the results presented in the previous chapter. Most notably, the argument that the greater the degree of male:female inequality in favour of men, the higher the risk for the entire reference group. The analysis of the data from the antenatal clinics found this very relationship.

While the other extensions to the theory cannot be easily tested, it is important to consider some of the possible implications. The following chapter considers the implications of the results outlined in the previous chapter and discusses the theory presented in this chapter.

Chapter 6

Conclusions

HIV in Africa is transmitted primarily through unprotected heterosexual sex. The sexual and health seeking behaviours of individuals play an important role in the determination of their risk of being infected. The discussions in this thesis add to the argument that, while these behaviours are important, it is the context in which they occur that is key. Contexts shape behaviour and the risk of that behaviour. This thesis has concentrated on one factor contributing to the risk context, the role of income.

Previous work in the area had highlighted a number of important issues, but there appeared to be a contradiction between different types of studies. Research on individual level risk found that the higher an individual's income the higher their risk of infection (Ainsworth and Semali, 1998), while country level studies found that the higher the income of a country the lower the prevalence (Over, 1998). Efforts to explain this apparent contradiction suggested that it had occurred because the individual level studies had measured the role of income at the early stages of the epidemic when knowledge of protective behaviours was not widespread (World Bank, 1997). This thesis has attempted to build on this explanation, discussing and examining how the relationship at a group and individual level, between income and risk may have changed as knowledge has become more widespread.

This development of existing work was achieved by applying the concepts of relative and absolute income to the question of what role income plays in the spread of HIV. The concepts of relative and absolute income, typically associated with the debate on differential health outcomes between wealthy countries, have allowed for a distinction to be drawn between risk associated with an individual's income independent of others and relative to other members of their reference group. This distinction facilitated the development of the theory of the dual role of income in the spread of HIV, which builds

on previous studies and explanations by examining how a change in income can increase and decrease risk, and discussing how the net effect is likely to be determined.

The Theory of the Dual Role of Income in the Spread of HIV hypothesises that two individuals with equal income, the same sexual preferences and of the same sex and age will have different levels of risk, depending on the average income of their reference group. The lower the average income of their group the higher the risk of their infection. This occurs for two reasons: firstly the lower the average income of the group, the higher the risk for everyone in that group and secondly, the lower the group average the higher their income is relative to others within the group.

The lower average income of the group is argued to affect group and individual risk as follows. Access to health services and education is lower at lower levels of income; as a result STIs are more prevalent and knowledge of protective behaviours more rare. The average risk of HIV infection among group members is, therefore, greater the lower the average income of the group. This affects the risk of an individual in the group in two ways. Firstly, their risk is affected by access to health services and knowledge; even if they themselves have greater knowledge, their behaviours will be constrained by what is deemed acceptable by the group. Secondly, the higher the average risk of the group, the greater the chance that sexual partners selected by an individual from the group are HIV positive. This increases an individual's probability of coming into contact with the virus and therefore their risk of infection.

The second risk factor associated with lower average group income is that the lower the average the higher an individual's income is relative to other members of the group. For men, this increases their power and attractiveness within the group, increasing, on average, their number of sexual partners, in turn increasing the risk that they will come into contact with HIV. For women, it increases the probability that they will partner with a male of higher relative income and associated higher risk.

Returning to the two identical individuals from different groups: if one individual were from a group with low average income, but at the upper end of the income distribution within that group, while the other was from the bottom end of the income distribution of a group with much higher average income, then the first individual would be at far greater risk of HIV infection than the second.

While an ideal data set to examine the explanatory power of this theory does not exist, there are other data sources the analysis of which produces results in line with expectations based on the theory. Income appears to increase the risk, at least for women, of individuals within groups with similar average incomes. But between groups with very different incomes, income appears to reduce the risk of the individuals within those groups. Furthermore, the data analysed in this thesis provided evidence of the importance of group characteristics in the determination of individual risk. The analysis suggested that within groups with high rates of female poverty and male:female income inequality the risk of all group members was increased. Similarly, the higher the rate of mobility of the group the greater the risk for the entire group. Finally, evidence was presented that suggests that the poor (reference groups with low average absolute income) living in wealthy areas[11], are worse affected than the poor in poor areas.

All of the results presented in this thesis support the argument that context and group characteristics are important in determining individual risk. These results are, therefore, in line with much of the literature reviewed in the first chapter of this thesis. Research has highlighted, among other things, the importance of STIs, access to health care, social constructs of acceptable behaviour, risks associated with poverty, the importance of information, and gender inequalities.

Taken collectively, the implications of these findings are substantial. Diverting, for a moment away from HIV, the implications for the prevention of HIV become clearer. Let us say, for the purpose of illustration, that my office is in a building on the side of a road

[11] In other words reference groups with low average absolute income are at greater risk if they live in close proximity to reference groups with high absolute income.

and the place where I eat lunch is on the other side of this road, directly opposite. Further along this road there are traffic lights. Each day I come out of my office to get lunch, but because I am over-worked I have very little time. So each day I come out, run across the road and buy lunch. Now I am not the only one doing this each day, a number of others in the same building do the same, because they too are over-worked. This is rather a dangerous exercise as the road is busy with cars and bicycles. Being hit by a bicycle won't kill you, but will slow you down and will therefore increase your chances of being hit by a car, which will kill you. The problem has got to such a stage that the manager of the building regularly calls us all in and tells us of the dangers of running across the road. But my lunch break is still too short, as I still have too much work, and the traffic light is still too far. I am sure that I will not get hit by a car, and while I could do without lunch and wait for dinner, really I'd rather not, so I still go out and run across the road, not every day now, but most of the time. The only difference when I do it now is that I know that it's not safe.

What is clear from the above example is that my behaviour is shaped by my circumstances and that the risk of that behaviour is determined by how busy the road is; in other words, by the context in which my behaviour occurs. While it is necessary for me to know how to protect myself, it is equally necessary for the context to be changed to allow me to use that knowledge. What then is the appropriate response? Less work and more time for lunch? Perhaps move the traffic lights, ban bicycles or build a bridge?

To prevent HIV, more than information will be needed. For this epidemic to be stopped, without a cure or a vaccine, will require fundamental changes in society. Targets for prevention efforts should be poverty, the distribution of wealth, factors promoting mobility, access to health care and other contextual factors that affect the spread of HIV.

References

Adler, M., Forster, S., Richens, J., and H. Slavin, 1996, 'Sexual Health and Care: Sexually Transmitted Infections: Guidelines for Prevention and Treatment,' *Health and Population Occasional Paper*. Overseas Development Administration, London.

Aggleton P., and K Rivers, 1999, 'Behavioural Interventions for Adolescents' in L Gibney, R., DiClemente and S. Vermund (eds.) *Preventing HIV infections in developing countries*, Plenum Publications, New York.

Ahlburg, D., and E. Jensen, 1998, 'The economics of the commercial sex work industry'. In Ainsworth, M., Fransen, L., and M. Over (eds.), *Confronting AIDS: Evidence from the developing world.* World Bank/EU, UK.

Ainsworth, M., and I. Semali, 1998, 'Who is most likely to die of AIDS? Socioeconomic correlates of adult deaths in Kagera region, Tanzania.' In Ainsworth, M., Fransen, L., and M. Over (eds.), *Confronting AIDS: Evidence from the developing world.* World Bank/EU, UK.

Ainsworth, M., Fransen, L., and M. Over (eds.), 1998, *Confronting AIDS: Evidence from the developing world.* World Bank/EU, UK.

Allen, S., Lindan, C., Serufilira, A., van de Perre, P., Rundle, A., Nsengumuremyi, F., Carael, M., Schwalbe, J., and S. Hulley, 1991,'Hunam Immunodeficiency Virus in Urban Rwanda: Demographic Behavioural Correlates in a Representative Sample of Chilbearing Women'. *Journal of American Medical Associ*ation 266 (12), 1 657-63.

Aneshensel, C., and C. Sucoff, 1996, 'The neighborhood context of adolescent mental health.' *Journal of Health and Social Behaviour.* 37, 293-310.

Auvert, B., Buve, E., Lagarde, E., Kahindo, M., Chege, J., Rutenberg, N., Musonda, R., Laourou, M., Akam, E., and H. Weiss, 2001, 'Male circumcision and HIV infection in four cities in Sub-Saharan Africa.' *AIDS* 15 (sup. 4), S31-S40.

Boulier, B., and R. Rosenzweig, 1984, 'Schooling, search and spouse selection: testing economic theories of marriage and household behaviour.' *The Journal of Political Economy*, 92(4), 712-732.

Campbell, C., 1997, 'Migrancy, masculine identities and AIDS: The psychosocial context of HIV transmission on the South African Gold Mines.' *Social Science and Medicine*, 45(2), 273-281.

Carael, M., 1995, 'Sexual Behaviour' In Cleland, J., and B. Ferry (eds.), *Sexual Behaviour and AIDS in the Developing World.* Taylor and Francis. Bristol, PA.

Case, A., and M. Yogo, 1999, 'Does school quality matter? Returns to education and the characteristics of schools in South Africa.' *NBER working paper*, National Bureau of Economic Research, Cambridge.

Cohen, D., 1993. 'Poverty and HIV/AIDS in sub-Saharan Africa.' *Issues paper No. 27.* UNDP, Geneva.

Cohen, D., 1998, 'Socio-economic causes and consequences of the HIV epidemic in Southern Africa: A case study of Namibia.' *Issues paper No. 31.* UNDP, Geneva.

Colvin, M., Gouws, E. and I. Kleinschmidt. 2000.The prevalence of HIV infection and associated risk factors in a South African working population. Unpublished.

Crane, J., 1991, 'Effects of neighborhoods on dropping out of School and teenage childbearing.' In: Jencks, C., and P. Peterson, (eds.), *The urban underclass.* The Bookings Institute, Washington, DC.

Deaton, A., 1999, 'Inequalities in income and inequalities in health.' *NBER Working Paper.* National Bureau of Economic Research, Cambridge.

Deaton, A., 2001a, 'Relative deprivation, inequality, and mortality.' *NBER Working Paper.* National Bureau of Economic Research, Cambridge.

Deaton, A., 2001b, *Health, inequality and economic development.* Working Paper, Princeton University, Princeton.

De Gruttola, V., Mayer, K., and C. Horsburgh, 1989, 'Infectiousness of HIV between male homosexual partners.' *Journal of Clinical Epidemiology* 42 (9), 849-56.

Deheneffe, J., Carael, M., and A. Noumbissi, 1998, 'Socioeconomic determinants of sexual behaviour and condom use.' In Ainsworth, M., Fransen, L., and M. Over (eds.) *Confronting AIDS: Evidence from the developing world.* World Bank/EU, UK.

Dorrington, R., and L. Johnson, 2001, *The ingredients and impact of the HIV/AIDS epidemic in South Africa and its provinces*, UNICEF, Johannesburg.

Eibner, C., and W. Evans, 2001, *Relative deprivation poor health habits and mortality.* Working Paper, Department of Economics, University of Maryland, Maryland.

Filmer, D., 1998, 'The socioeconomic correlates of sexual behaviour: A summary of results from an analysis of DHS data.' In Ainsworth, M., Fransen, L., and M. Over, (eds.). *Confronting AIDS: Evidence from the developing world.* World Bank/EU, UK.

Gersovitz, M., Jacoby. H., Dedy, S., Tape, A. 1995, '*The balance of self reported Hetrosexual Activity in KAP Surveys and the AIDS Epidemic in Africa*'. Mimeo, Department of Economics, Johns Hopkins University, Baltimore, MD.

Greene, W. 1997. Econometric Analysis 3rd Addition. New Jersey. Prentice Hall.

Haverskos, H., and R. Battjes, 1992, 'Female to male Transmission of HIV.' *Journal of the American Medical Association*, 268 (14), 1855-56.

Hunt, C., 1989, 'Migrant Labour and Sexually transmitted diseases: AIDS in Africa.' *Journal of Health and Social Behaviour*, 30(4), 353-373.

Kennedy, P. 1998. A Guide to Econometrics. Cambridge. MIT press.

Kirunga, C., and J. Ntozi, 1997, 'Socio-economic Determinants of HIV Serostatus: A study of Rakai District, Uganda'. *Health Transition Review*, Supplement to Volume 7: Vulnerability to STD/HIV Infection and effects of AIDS in Sub-Saharan Africa, 175-88.

Lurie, M., 2000a, 'Migration and AIDS in southern Africa: a review.' *South African Journal of Science*. 96, 343 – 347.

Lurie M., William B., Sturm A., Garnett G., Mkaya-Mwamburi D., and S.S. Abdool Karim, 2000, 'HIV discordance among migrant and non-migrant couples in South Africa.' *Paper presented at the XIII AIDS Conference*, Durban.

Lynch, J., Smith, G., Kaplan, G., and J. House, 2000, 'Income inequality and mortality: importance to health of individual income, psychosocial environment or material conditions' *British Medical Journal*, 320, 1200-1204.

Mastro, T., and I. de Vincenzi, 1996, 'Probabilities of Sexual HIV-1 Transmission.' *AIDS*, 10 (suppl. A), S72-82

Meursing K., Sibindi F. 1995. `Condoms, family planning and living with HIV in Zimbabwe', Reproductive Health Matters, 5, 56-67.

Miller, D., and C. Paxson, 2000, *Relative income race and mortality*. Working Paper, Princeton University, Princeton.

Moses, S., Bradley, J.E., Nagelkerks, N.J., and A.R. Ronald, 1990, 'Geographical patterns of male circumcision practices in Africa: association with HIV seroprevalence.' *International Journal of Epidemiology*, 19, 693 – 697.

Msamanga, G., Urassa, E., Spiegelman, D., Hertzmark, E., Kapiga, S., Hunter, D., and W Fawzi. 1996, `Socioeconomic status and prevalence of HIV infection among pregnant women in Dar es Salaam, Tanzania'. Abstract TuC2464, *Eleventh International AIDS Conference*. Vancouver, B.C., Canada, July 7-12.

Over, M., 1998, 'The effects of societal variables on urban rates of HIV infection in developing countries: An exploratory analysis.' In Ainsworth, M., Fransen, L., and M. Over (eds.). *Confronting AIDS: Evidence from the developing world*. World Bank/EU, UK.

Preston, S., 1975, 'The changing relationship between mortality and the level of economic development.' *Population Studies*, 29, 231-48.

Rivers, K., and P. Aggleton, 1999, *Men and the HIV Epidemic*. Thomas Coram Research Unit Institute of Education, University of London, London.

Ross, C., 2000, 'Walking, exercising, and smoking: does neighborhood matter?' *Social Science and Medicine,* 51, 265-274.

Stattistics South Africa. 1996 South African National census. StattSA.

Stillwaggon, E., 2000, 'HIV transmission in Latin America: comparison with Africa and policy implications'. *South African Journal of Economics*, 68(5), 985-1011.

Stover, J., 1999, *AIM: A computer program for making HIV/AIDS projections and examining the social and economic impacts of AIDS*. The Futures Group International, Washington.

UNAIDS. 2000, *Report on the global HIV/AIDS Epidemic*. UNAIDS, Geneva.

UNAIDS/WHO. 2000, *Consultation on STD interventions for preventing HIV: what is the evidence?*. UNAIDS, Geneva.

Whiteside, A., 2002, 'Poverty and HIV/AIDS in Africa.' *Third World Quartely*, 23(2), 313-332.

Whiteside, A., and C. Sunter, 2000, *AIDS: The challenge for South Africa*. Human and Roussea Tafelberg, Cape Town.

Wilkinson, R., 1996, *Unhealthy Societies: The affliction of inequality*. Routledge, London and New York.

Wilkinson, R., 1997, 'Health Inequalities : Relative or absolute material standards?' *British Medical Journal,* 314, 591-595.

Williams, B., Gilgen, D., Campbell, C., Taljaard, D., and C. MacPhail, 2000, *The natural history of HIV/AIDS in South Africa: A biomedical and social survey in Carletonville.* Council for Scientific and Industrial Research. Johannesburg.

World Bank. 1998, *Confronting AIDS*, Oxford University Press, New York.

World Health Organisation. 1994, *Women and AIDS: Agenda for Action*. WHO, Geneva.

Zwi, A., and J. Cabral, 1991, 'Identifying high risk situations for preventing AIDS.' *British Medical Journal*, 303, 1527-1529.

Appendix

Table A1: Variables examined but excluded[12] from the final model

Variable	Mean	Std. Dev	Description
Total population	8126.64	5265.982	Total population of the magisterial district in which the clinic was located
Age at first birth	19.44	4.14	Average age of women in the magisterial district at the time of the birth of their first child
National roads	0.56	0.50	Dummy variable indicating if the magisterial district had a national road passing through it
Urbanisation	0.44	0.36	Proportion of households in the magisterial district classified as urban
Male average income	436.88	209.47	Average income of all adult males in the magisterial district
Female average income	234.14	132.28	Average income of all adult females in the magisterial district
Working male average income	1225.03	282.78	Average income of all employed men in the magisterial district
Working female average income	880.07	234.62	Average income of all employed women in the magisterial district
Female to male income	0.53	0.14	Female average income as a proportion of male average income
Working female	0.72	0.07	Working female's income as a

[12] Variables were excluded if they were insignificant when included only with age variables and with the final model.

income to working male income			proportion of working men's income
Proportion of households in informal housing	0.08	0.10	Households living in informal housing as a proportion of all households in the magisterial district
Not working	0.71	0.07	Proportion of adults not employed
Visitors	0.02	0.01	Proportion of total adult population classified as visitors on census night
Long distance work related mobility	0.05	0.04	Proportion of working adults working outside the province
Short distance work related mobility	0.04	0.05	Proportion of working adults working outside the magisterial district of the household but within the same province
Long distance male work related mobility	0.08	0.08	Proportion of working men working outside the province