

Chapter 1

Introduction

Micro health insurance protects the poor from the financial hardship of health expenditures but suffers from low demand. This thesis analyzes three potential reasons for low take-up and this chapter embeds the thesis in the literature. First, according to prospect theory, consumers derive value not from consumption itself but from changes relative to a reference point. Theoretical predictions in Chapter 2 suggest that reference-dependent preferences indeed reduce demand for insurance in a context of perfect financial markets, but may increase demand when consumers face liquidity constraints. A second hypothesis is that there is adverse selection, meaning that households are more likely to enroll high-risk members than low-risk members. However, based on an insurance program in rural Nigeria, Chapter 3 finds no evidence of adverse selection within households. Finally, Chapter 4 hypothesizes that potential clients forgo insurance because they rely on assistance from their social network. A framed field experiment provides evidence of this theory. Microfinance clients opt out of individual insurance to free-ride on contributions from fellow group members. Insurance offered at the group level binds them to the group optimum and enhances demand.

1.1 Background

Health is one of the most pressing challenges for households in poverty. Without access to prepaid sources of health finance, the poor need to pay for health care directly out-of-pocket. When households cannot pay these out-of-pocket health expenses, they go to great lengths to raise the money or postpone seeking health care until symptoms progress into a medical as well as financial emergency. Every year an estimated 150 million people suffer financial hardship because they are forced to spend a large share of their incomes on health services out-of-pocket, pushing many of them into extreme poverty (Xu et al., 2007).

Health risk has tremendous welfare effects both ex-post - once a person becomes ill or injured - and ex-ante - when trying to mitigate the risk. Numerous studies show that households deal with the financial hardship by resorting to ex-post coping strategies such as borrowing, selling productive assets, reducing consumption, working longer hours or taking their children out of school (Foster, 1994; Jacoby and Skoufias, 1997; Gertler and Gruber, 2002; Leive and Xu, 2008). These are costly solutions that disrupt consumption plans, absorb households' assets and hamper

the development of human capital.

The mere presence of uninsured risk can reduce investments in productive activities even before a shock occurs. Households protect themselves from adverse shocks ex-ante by smoothing their incomes. For instance, they make conservative production or employment choices and diversify economic activities (Morduch, 1995), or save with precautionary motives in low-return buffer stocks (Deaton, 1991). If households could insure themselves against the risk, they would not need to engage in these low-return activities. Elbers et al. (2007) show that also through this channel, risk can severely impede wealth accumulation and economic growth.

Health risk has significant adverse effects in low-income countries for at least three reasons. First, low-income countries have limited resources to invest in health care. According to the ‘first law of health economics’ (Van der Gaag and Stimac, 2012), health spending is log-linearly increasing with a country’s GDP. A low GDP hence means that very few resources are available for health. Even fewer resources are used to prepay health care, for instance through public spending and health insurance. Instead, low-income countries rely heavily on out-of-pocket payments with ratios of prepaid to out-of-pocket sources as low as 20:80. As a comparison, this is 80:20 in high-income countries (Preker et al., 2007).

Second, direct payments can create financial hardship because of limited access to formal financial instruments like savings accounts, credit and insurance. Half of the world is unbanked (Chaia et al., 2012), in part due to financial market imperfections such as information asymmetries and transactions costs. Hidden action or moral hazard can reduce repayment rates, and this makes banks hesitant to lend without collateral (Karlan and Zinman, 2009). Documenting credit histories helps limit information asymmetries. De Janvry et al. (2010) show this for the case of Guatemala, where a major microfinance lender set up a credit bureau to exchange information about clients’ repayment performance. However, in most low-income countries, weak institutions, limited property rights, as well as a lack of collateral and credit histories restrict access to financial markets (Beck et al., 2009).

Third, less formal insurance strategies exist, but these have severe shortcomings. Microfinance institutions (MFIs) for instance increase access to credit, building on social collateral to service the poor. Nevertheless, randomized studies on the impact of microfinance often do not find an increase in health spending (Banerjee et al., 2010a; Karlan and Zinman, 2010). Emergency loans - for instance for health expenditures - are typically considered too risky because of low repayment rates, and MFIs offer loans mainly for productive investments. They do so through group-based lending, meaning that clients are jointly liable for defaulting group members. This incentivizes clients to do ‘the MFI’s work’: screening and monitoring their fellow group members, which will limit adverse selection and moral hazard, so that loans are invested productively and repaid (Armendariz and Morduch, 2010).

Alternative informal insurance strategies do exist, but again with severe shortcomings. In case of financial emergencies, households typically receive gifts and loans from their social network, for instance from neighbors or the extended family. But within these risk-sharing networks, shocks affect consumption substantially, suggesting that the risk is not fully shared (e.g. Cochrane, 1991; Townsend, 1994; Udry, 1994; Dercon and Krishnan, 2000). Thus, informal insurance arrangements provide incomplete insurance. This is the case for covariate risks such as droughts and floods, but also for idiosyncratic risks like illnesses and injuries.

To summarize, health risk can tremendously reduce welfare. Low-income countries have limited resources for public health care. Without the proper financial instruments to cope with private health expenditures, households resort to costly ex-post coping strategies. Anticipating this, households may engage in low-return activities to smooth their incomes and consumption, limiting wealth accumulation even further. Health insurance as a pre-paid health financing tool may help protect the poor from this financial hardship.

1.2 Micro health insurance

Policy-makers are nowadays expanding ‘micro health insurance’ to protect the poor from catastrophic health expenditures. A typical micro health insurance product is an affordable health plan for the poor with basic benefit packages that are easy to understand. A wide range of different micro health insurance products exists. Most schemes are partly subsidized and aim at increasing private willingness to pre-pay for health care, so that the subsidies can be phased out. To limit information asymmetries like adverse selection and moral hazard, many programs restrict benefits to mainly inpatient health care (i.e. hospitalizations). They collaborate with pre-existing institutions like cooperatives and microfinance banks to minimize transaction costs. Publicly funded programs can be heavily subsidized and population- or community-based. High subsidies allow these schemes to offer more comprehensive benefits, including outpatient care.

Micro health insurance started as a promising health financing tool, sometimes called the new ‘microfinance revolution’. Early impact evaluations suggested that these programs increase health care utilization and provide financial protection by reducing out-of-pocket spending (Ekman, 2004). Recent impact evaluations have expressed more skeptical views. In Acharya et al. (2012), for instance, only a small majority of reviewed studies find a positive impact on health care utilization and insurance does often not reduce (and sometimes even increase) average out-of-pocket payments. A few studies find a reduction in catastrophic out-of-pocket payments.¹

Low enrollment and renewal rates are among the key challenges for these programs. Many

¹A reduction in catastrophic out-of-pocket payments might be a better indicator for welfare improvements than a reduction in average out-of-pocket payments. An increase in health expenditures may reflect increased access to health care and is not a bad impact per se.

micro health insurance schemes seem very attractive at first glance, but take-up can remain as low as for instance 5.2 percent in a program in Burkina Faso (Gnawali et al., 2009). Although large subsidies and discounts enhance demand, the majority of potential beneficiaries remain uninsured even at heavily discounted premiums (Thornton et al., 2010; Dercon et al., 2012; Polimeni and Levine, 2012). This does not only harm the sustainability of microinsurance but also reduces the power to detect an impact, and might hence explain why evaluations do not find more positive results.

Low take-up rates are inconsistent with basic economic frameworks for health insurance demand. These models predict that any risk averse consumer purchases actuarially fair insurance. To illustrate this, consider Figure 1.1. This figure shows utility as a function of consumption. The consumer has an initial endowment of wealth w . Consuming this endowment yields utility $U(w)$. But with probability p , the consumer needs to spend her endowment on health care, leaving her with zero consumption and zero utility. Expected utility is a linear combination of the two utility levels, $(1 - p)U(w)$, indicated by ‘Uninsured’ on the vertical axis.

When the consumer purchases insurance at an actuarially fair premium wp , earnings are $w(1 - p)$ with certainty. The associated utility is indicated by ‘Insured’. Since the utility function is concave, i.e. the consumer is risk averse, utility with insurance exceeds expected utility without insurance. Thus, under the regular assumption of risk aversion, it is optimal to buy actuarially fair insurance. Low take-up of subsidized micro health insurance contradicts this principle (Mossin, 1968; Doherty and Schlesinger, 1983).

1.3 Why low take-up? Reconciling theory and evidence

The model in Figure 1.1 is very basic and abstracts from a number of mechanisms that could result in low take-up. This section reviews the literature on four potential mechanisms that Figure 1.1 does not incorporate: institutional features characterizing low-income markets, non-standard preferences, beliefs and the decision-making process. The review starts from the assumption that taking insurance in principle enhances welfare, so that increased take-up is desirable. Because intentions do not necessarily result in take-up, I focus on revealed demand and do not discuss studies on stated willingness to pay (e.g. Gustafsson-Wright et al., 2009). Further, since insurance is an investment in health, the discussion will go beyond take-up of micro health insurance, reviewing a more general literature on the adoption of preventive health technologies in developing countries (Kremer and Glennerster, 2011). This literature includes a larger number of studies that test which mechanisms drive adoption. For reviews focusing exclusively on the demand for micro (health) insurance, see Schneider (2004) and Matul et al. (2013).

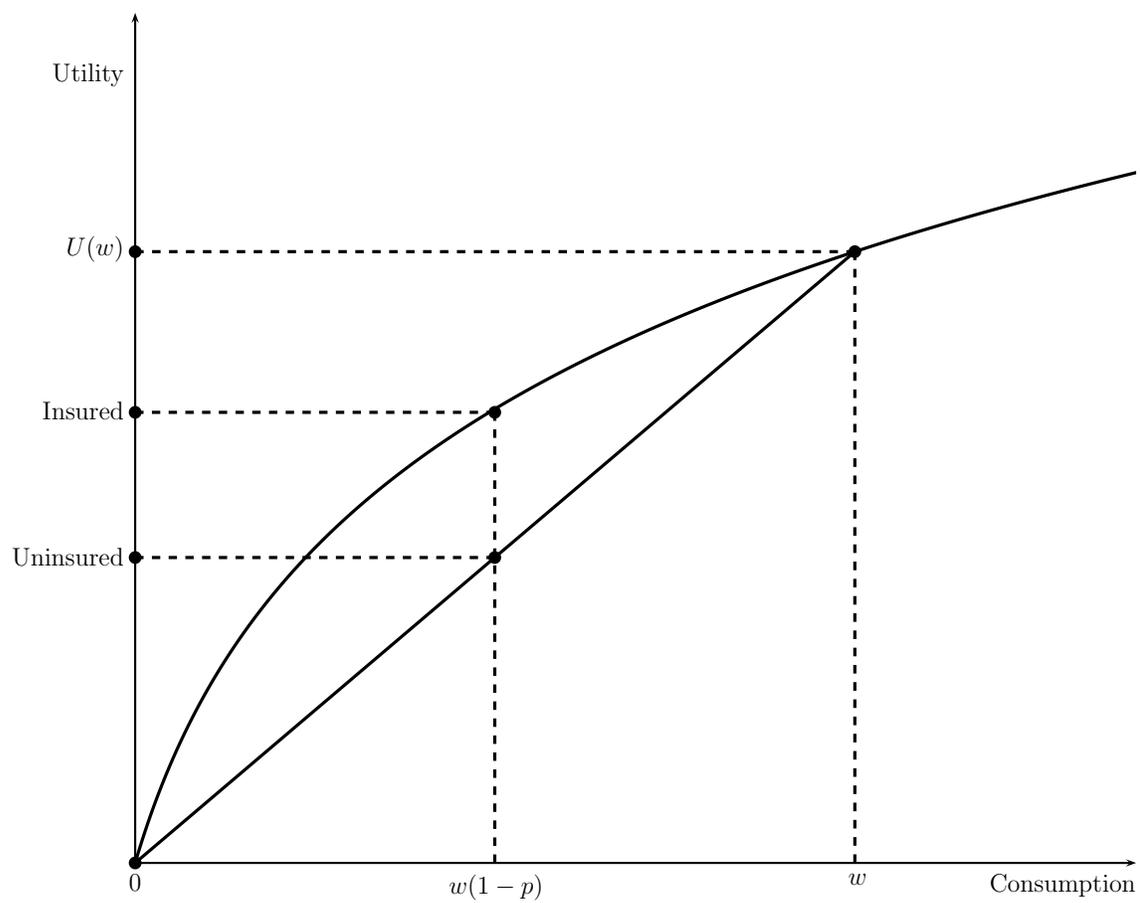


Figure 1.1: Risk aversion and demand for insurance

The institutional environment

A first class of mechanisms relates to the institutional environment of micro health insurance. Low take-up cannot be seen in isolation from the context in which insurance operates. Liquidity considerations, high transaction costs, background risk, and a low quality of supplied health care - typical struggles for households in low-income countries - inevitably influence demand.

To start, in imperfect financial markets, the cost of credit is high. This may harm households' ability to pay the annual insurance premium - not because they do not have the money, but because they have insufficient funds on hand to pay the premium in one go and are unwilling to borrow at high interest rates.² Investments in preventive health and microinsurance are generally consistent with this liquidity interpretation. Field experiments show that households who just received cash are more likely to purchase rainfall insurance (Cole et al., 2013) and mosquito nets (Hoffmann et al., 2009) compared to households that have to pay from their own pocket. Also microfinance loans, which offered at more generous interest rates than loans from informal money lenders, increase the demand for mosquito nets (Tarozzi et al., 2011).

An important caveat is that micro health insurance - open for enrollment throughout the year - is qualitatively different from rainfall insurance, where payments are typically due early in the season when cash is low. In addition, these studies potentially do not measure liquidity considerations but the effect of participating in the experiment, so-called experimenter demand effects. Participants who have just received cash or a loan might feel a social obligation to purchase the offered technology. They might also perceive the cash or loan as a windfall. Since they did not decide on how to spend the money yet, allocating it to new technologies will be easier.

A number of field experiments randomize the micro health insurance premium (Thornton et al., 2010; Dercon et al., 2012; Polimeni and Levine, 2012). In general, they show that discounts increase take-up significantly, especially among the poor for whom access to credit is restricted most. This is again consistent with the hypothesis that liquidity considerations are a barrier to enrollment but does not prove it. Models with perfect financial markets also predict price sensitivity, and even more so for the poor with low consumption and very tight budget constraints. Liquidity offers a plausible explanation, but the evidence is not as convincing as postulated in earlier literature reviews, e.g. Matul et al. (2013).

Transaction costs are a second potential barrier to enrollment. Convenience and distance are important determinants of the demand for health goods. In Thornton (2008), randomized distance to test centers reduces pick-up rates of HIV test results. In Kremer et al. (2011), individuals are willing to walk only 3.5 minutes further to collect water from a spring that is protected from

²Such liquidity considerations might however also *increase* the willingness to pay for insurance since credit is an alternative risk-coping tool and potentially a substitute for insurance (Gollier, 2003).

ground contamination. Thornton et al. (2010) randomize the location where market vendors could register for micro health insurance. Most clients had to enroll at a central office, where registration took about one day's time. Some clients were offered the opportunity to register at their market booth. This convenience increased take-up by more than half the effect of a 6-month subsidy.

However, transaction costs are not easily reduced without large investments in institutions, infrastructure or operational costs. Registering every client individually at her market booth will increase operational costs and hence also the premium. This might not be cost-effective in the long run. An alternative is to offer small incentives that motivate consumers to register and renew despite high transaction costs. These nudges can be small and cost-effective, as demonstrated in Banerjee et al. (2010b). This study in India found that very small - but appreciated - incentives induced parents to travel up to five kilometers to get their children immunized.

A third explanation for low take-up is that potential clients face a multitude of other risks, from ethnic conflicts to agricultural disaster and from death in the family to volatile commodity prices. Such uninsurable background risk may affect optimal demand for insurance depending on the correlation between the uninsurable losses and the health risk (Doherty and Schlesinger, 1983). Specifically, if health expenditures and other uninsurable risks correlate negatively, other risks serve as a hedge for the health risk. This reduces demand for insurance. And when liquidity is a major consideration, enrolling might be suboptimal even when the risks are uncorrelated, since insurance earmarks scarce resources for health rather than for precautionary savings to cope with risk in general. To my knowledge, no study investigates how background risk affects investments in preventive health or micro health insurance.

A fourth concern is the poor quality of health care provided by insurance, as a number of qualitative focus group discussions suggest (Criel and Waelkens, 2003; De Allegri et al., 2006; Basaza et al., 2008). Also quantitative studies show that being dissatisfied with the quality of supplied health care is associated with lower enrollment (Platteau and Ugarte, 2013) or renewal (Dong et al., 2009). This is not surprising given that many poor countries have few resources for health systems. In response to the relatively low quality of health care that insurance schemes provide, the poor may turn back to fee-for-service care and forgo health insurance. Improving the quality of supplied health care should hence be an integral component of any micro health insurance intervention.

Nevertheless, these four institutional features of low-income countries do not provide the full answer to the low take-up puzzle. Demand remains lower than expected even in programs that upgrade partnering health facilities or enroll clients at their homes to reduce transaction costs. Heavily discounted premiums, which should remove barriers associated with liquidity and background risk, increase demand but not to sustainable levels. Hence, other - behavioral -

factors likely hamper take-up as well. The next three sections discuss these factors.

Preferences

A first class of behavioral models relates to time and risk preferences. Figure 1.1 assumes that the timing of the insurance purchase, the realization of the health risk, and consumption is irrelevant. In reality, the premium payment and the realization of risk affect consumption at different points in time. As a result, impatient consumers who do not sufficiently value future insurance benefits do not take insurance. In addition, the figure assumes that consumers are strictly risk averse. For risk-seeking consumers with convex utility, insurance will not be utility-enhancing.

Risk-loving preferences and impatience are however difficult to reconcile with the stylized fact in Figure 1.2 that demand for health prevention technologies is very price-elastic around zero (Kremer and Glennerster, 2011). This applies to different goods in various countries; from deworming pills and mosquito nets in Kenya to water disinfectant in Zambia and soap in India. Individuals in different contexts appear indifferent as to whether they will invest or not at a price just above zero. In traditional economic models, this finding is implausible, even when including the institutional features like liquidity constraints described above.

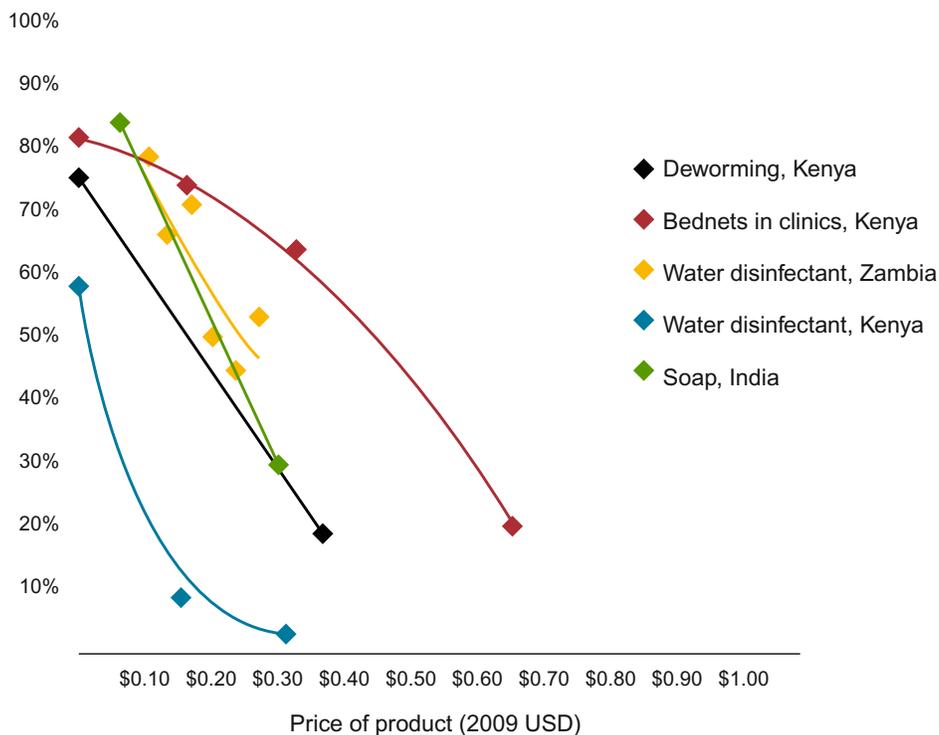


Figure 1.2: Demand for health prevention (Source: www.povertyactionlab.org)

Hyperbolic discounting is one explanation for the stylized fact that demand is highly price-elastic around zero. Hyperbolic discounters discount at a higher rate between the present and

immediate future than between two future time periods (Loewenstein and Prelec, 1992). As a result, they prefer \$100 today to \$110 tomorrow while preferring \$110 in a year and 1 day to \$100 in a year. Due to this present bias, they cannot commit themselves to invest (Laibson, 1997). Even a small cost associated with enrolling in insurance (e.g. a heavily discounted premium, or the time needed to register) can then have a large effect on take-up. A naive hyperbolic discounter may plan to enroll in the future, but once the future becomes near, these small costs are perceived too high. This potentially explains the demand pattern in Figure 1.2.

Consumers with the intention to purchase insurance in the future may realize that they are not committed to pay the insurance premium once they have to enroll or renew. A sophisticated consumer, aware of this self-control problem, will seek ways to commit herself (Strotz, 1955; Pollak, 1968); for instance by earmarking savings to purchase insurance in the future. Consistent with this hypothesis, accounts that returned savings conditional on passing a urine test helped to quit smoking (Giné et al., 2010); commitment savings devices such as safe boxes and health savings accounts increased investments in preventive health (Dupas and Robinson, 2013); and in an experiment in India, about half of microfinance clients committed themselves in advance to re-treat their insecticide-treated bed nets (Tarozzi et al., 2011). By contrast, Ito and Kono (2010) do not find that hyperbolic discounters are less likely to enroll in health insurance, measuring present bias directly through incentivized games. However, trust in the surveyor easily confounds such measures.³

Whether commitment devices help increase take-up through this channel has not been tested yet. When designing studies to this end, it is important to realize that commitment devices will not work for all. Demand for commitment will only exist if hyperbolic discounters are sophisticated enough to recognize their self-control problem. Less sophisticated consumers believe they can commit to their plans and do not use commitment devices, but then fail to carry through their plans. Moreover, a delicate trade-off exists between commitment and flexibility. The typical microinsurance client lives - as discussed earlier - in an environment of imperfect credit markets and background risk. Being unable to predict their future state, consumers may be unwilling to commit themselves too far in advance.

Further, ex-post, when struck by uninsurable losses, an insured client may regret having committed her resources to health. Take the anecdotal example of a health insurance program for dairy farmers in Kenya. This program partners with cooperatives partly because they can automatically deduct the premium from farmers' dairy sales. The automatic payment can be interpreted as a commitment device. Farmers however appear to escape the commitment by side-selling their dairy. They do so for instance when cash is low, which is the case especially

³In the example above, also a lack of trust in the experimenter can explain why an experimental subject prefers \$100 today to \$110 tomorrow while the reverse holds in a year from now. See Andreoni and Sprenger (2012) for a discussion on how to limit a confounding effect of trust.

in the dry season. Side-selling may harm collective welfare through its adverse effects for the cooperative.

The success of commitment devices hence hinges on both preferences and the institutional context, such as a lack of affordable credit and the presence of background risk. At this stage, we know little about how these two interact.

Beliefs

A second class of behavioral mechanisms focuses on potential clients' beliefs about insurance and the probability that they fall ill. In this context, I discuss trust, understanding, social learning and awareness of the health risk.

First, households may lack trust in health insurance providers. When the insurance provider is likely to default, forgoing actuarially fair insurance can be optimal. Doherty and Schlesinger (1990) show that if total contract nonperformance occurs with some positive probability, the relation between demand for insurance and risk aversion has an inverse-U shape. For less risk averse consumers, demand is increasing in risk aversion, as the basic model in Figure 1.1 would predict; but for more risk averse consumers, demand is decreasing in risk aversion. Very risk averse clients who do not trust the insurance provider will opt out because they strongly dislike the worst-case scenario in which they pay for insurance, have a claim, but do not receive any benefits from the insurance provider.

Dercon et al. (2011), studying health insurance decisions among Kenyan tea farmers, provide empirical evidence of this theory. Demand is increasing in risk aversion for less risk averse farmers, but decreasing in risk aversion for more risk averse farmers, and demand is low in particular among farmers who do not trust the insurance provider. Solving this lack of trust is not straightforward though. Dercon et al. (2012) experimented with peer referrals to improve trust and enhance take-up. This actually had a counterproductive effect on demand, as it created associations with notorious pyramid schemes. Again, the institutional context appears an important consideration when designing micro health insurance programs.

Potential buyers may also not fully understand insurance. Platteau and Ugarte (2013) for instance find that survey respondents who understood what insurance is (a redistribution of resources from lucky to unlucky individuals, as they call it), are more likely to enroll in insurance. If potential clients rationally process information, then insurance education should enhance enrollment when conveying new information that the client did not have before. Three randomized controlled trials (henceforth 'RCTs') test this hypothesis for microinsurance. Cai (2012) finds that insurance education increases knowledge as well as demand for agricultural insurance. By contrast, Dercon et al. (2012) and Bonan et al. (2012) find that insurance literacy increases understanding but not demand for health insurance. Thus, part of the low take-up puzzle is why improved understanding does not translate into increased demand.

The effect of insurance education and learning in general will depend on the social network. Peer effects can move in two directions though. On the one hand, studies find a positive association between the probability of adopting a preventive technology and the number of social links to people who received bed nets (Dupas, 2010) or menstrual cups (Oster and Thornton, 2012). Also for agricultural microinsurance, Cai (2012) finds that social learning enhances take-up. On the other hand, Kremer and Miguel (2007) find negative peer effects on the adoption of deworming drugs, especially for families with more schooling. They argue that more educated families start out with favorable beliefs about the technology but revise their beliefs downwards as they acquire more information. Social learning may hence affect take-up of micro health insurance in two ways, depending on the message conveyed within the network.

Finally, awareness about private health risks is low. A systematic review of hypertension in sub-Saharan Africa for instance finds that less than 40 percent of hypertensive adults were aware of their high blood pressure (Addo et al., 2007). Not surprisingly then, the literature yields mixed findings on whether take-up is higher among individuals with higher risks. In rural Cambodia, households who select into insurance at high prices have significantly higher health expenditures than those selecting into insurance at low prices (Polimeni and Levine, 2012). This suggests that health risk influences the willingness to pay. In rural China, households with sick members are more likely to purchase insurance (Wang et al., 2006; Zhang and Wang, 2008). However, Parmar et al. (2012) and Dercon et al. (2012) do not replicate this finding in Burkina Faso and Kenya, respectively. In other words, the literature does not provide conclusive evidence of adverse selection.

The decision-making process

A final class of models deviates from standard theories for decision-making under risk. The model above assumes that clients optimize their expected utility - the sum over all possible utility outcomes multiplied by their probability, which is $U(w(1 - p))$ with insurance and $(1 - p)U(w)$ without insurance in the example of Figure 1.1. But expected utility does not always drive decisions. This review discusses two alternative decision models - limited attention and prospect theory, in which preferences are reference-dependent.

First, potential clients may not buy health insurance until the health risk becomes salient to them, for instance after a household member experiences a health shock. Wang et al. (2006) and Zhang and Wang (2008) observe this type of adverse selection, which could suggest that limited attention drives low take-up. To test this assumption more rigorously, one could use the prediction that information about individual health risk factors, e.g. from a check-up, *increases* the willingness to pay if a person learns that she is at an increased risk, but does *not decrease* the willingness to pay if the person learns that she is healthy. Several papers reviewed by Kremer and Glennerster (2011) show that the adoption of health prevention technologies responds

asymmetrically to such information, but it has not been tested for micro health insurance take-up.

Second, consumers may not buy health insurance if they frame health expenditures as a big loss and have convex utility over losses. This is a key idea in prospect theory, which starts from the assumption that decisions depend on the frame used. Consider the following example (Kahneman and Tversky, 1979):

(i) In addition to whatever you own, you have been given 1,000. You are now asked to choose between (A) 1,000 with probability 0.5, or (B) 500 for sure.

(ii) In addition to whatever you own, you have been given 2,000. You are now asked to choose between (A) -1,000 with probability 0.5, or (B) -500 for sure.

The majority of participants chooses option B in the first problem and option A in the second. When viewed in terms of final states, the two problems are however identical. In general, participants are more risk averse when decisions are framed as a sure gain than when framed as a sure loss (Wakker, 2003). This property, called diminishing sensitivity, explains why participants prefer B over A in the first problem and reverse their preferences in the second problem. In addition, prospect theory assumes that losses hurt more than gains are pleasant in mixed gambles with both losses and gains - a property called loss aversion.

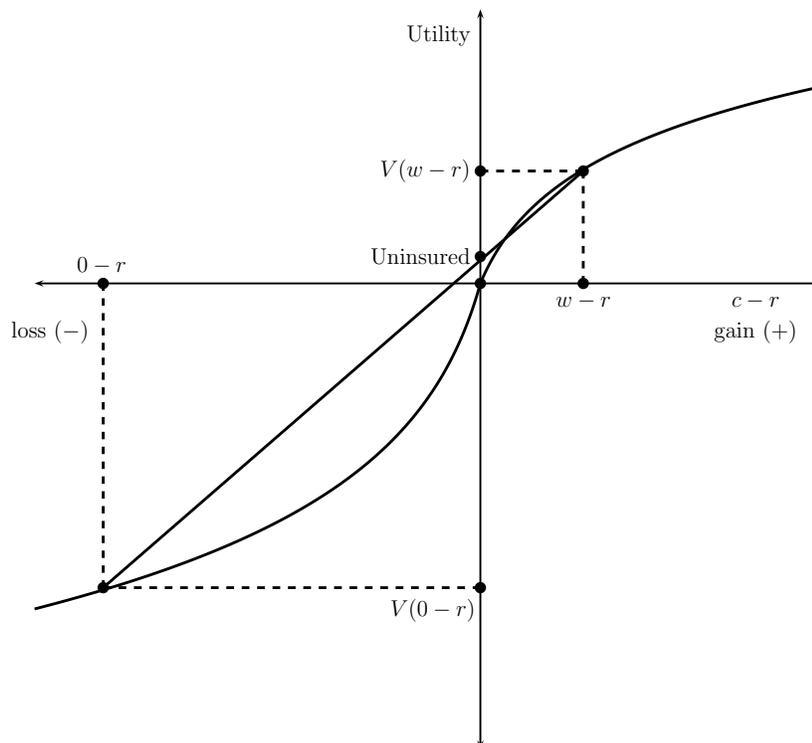


Figure 1.3: Prospect theory and demand for insurance

Figure 1.3 illustrates the insurance decision given these preferences. Consumers derive value from changes in consumption c relative to the reference point r . Loss aversion results in a kink

at the origin, where consumption equals the reference point, and by diminishing sensitivity, consumers are risk averse only towards gains and risk-seeking towards losses. This results in concave and convex utility in the gains and loss frame, respectively. In other words, gaining an additional \$10 on top of \$1,000 will be less pleasant than gaining \$10 on top of \$100. Likewise, losing an additional \$10 on top of \$1,000 will hurt less than losing \$10 on top of \$100.

Suppose an insured consumer's wealth, after paying the premium, is at the reference point, i.e. she does not gain or lose. Her utility is zero. An uninsured consumer is healthy with probability $1 - p$, yielding $V(w - r)$, and ill with probability p , yielding $V(0 - r)$. Utility without insurance, labeled 'Uninsured', now exceeds utility with insurance, which is zero. It is optimal not to purchase insurance.⁴ Because preferences satisfy diminishing sensitivity, demand is potentially larger if insurance covers smaller, more frequent losses. Anecdotes from practitioners introducing microinsurance often support this prediction.

The theory's main limitation is the infinite degrees of freedom associated with an exogenous reference point. A number of applications of prospect theory, starting from Kőszegi and Rabin (2006), bridge this gap by treating the reference point as an endogenous belief held prior to consumption. Changes in beliefs then yield gain-loss utility. Recent experiments (Ericson and Fuster, 2011; Abeler et al., 2011) and empirical studies (Crawford and Meng, 2011; Bronchetti et al., 2011) support the hypothesis that consumers compare outcomes to their plans or expectations. An open research question is how this preference structure affects the adoption of preventive health technologies and micro health insurance.

What is missing?

To summarize, enrollment rates in micro health insurance are lower than expected and demand is highly price-sensitive. This can be explained partly by features of the institutional environment, and partly by behavioral factors like non-standard preferences, beliefs and the decision-making process. Interventions aimed at removing these barriers enhance enrollment but do not reach sustainable levels. Further, different mechanisms appear at work in different settings. The remainder of this section identifies three gaps in the existing literature and discusses how these are addressed by the three core chapters in this thesis.

A first gap is the lacking knowledge of how institutional and behavioral mechanisms interact. The literature on institutional constraints is fairly separate from behavioral decision-making theories. The two will however be interrelated. For instance, scarcity affects how people with limited attention take decisions. Shah et al. (2012) show that scarcity induces participants in a laboratory experiment to direct their attention towards immediate needs. Outside the laboratory,

⁴Note that this result does not always hold under prospect theory. It requires sufficiently small loss aversion and strong risk-seeking for losses. Further, the figure ignores probability weighting, an essential feature of prospect theory. Consumers overweight high-impact low-probability events and underweight low-impact high-probability events. Consumers would hence overweight the health expenditures, making them more risk averse.

this could imply limited attention for future health expenditures, explaining low take-up of health insurance among the poor. Prospect theory potentially interacts with institutional features as well, but a major question is how.

Second, most non-experimental studies report partial correlations and cannot infer causality. Without exogenous variance in the explanatory variables of interest, these studies do not shed light on the channel behind the correlation. Policy conclusions from causal evidence are much stronger. A core part of economic science is hence to deal with confounding unobserved heterogeneity, for instance by using panel data. Unfortunately, the majority of studies on the demand for micro health insurance is cross-sectional, with one observation for every decision-maker, and often unable to deal with unobserved heterogeneity. Panel data will greatly enhance the quality of empirical evidence on solutions to the low take-up puzzle.

Third, most experimental studies on micro health insurance test whether a particular intervention enhances demand, for instance whether discounts, financial literacy training or on-site registration increase take-up. These studies however do not identify through which mechanisms the interventions have an effect, e.g. through liquidity considerations, social learning, self-control problems, etc. These so-called ‘black-box’ answers limit the external validity of RCTs (Deaton, 2010; Harrison, 2011). Alternative methods like laboratory experiments offer a higher degree of control, which helps identify which mechanisms make a successful intervention. Laboratory experiments range from experiments with university students (‘conventional’ experiments) to ‘framed field experiments’ that frame the experimental task in a field context. The latter type enhances the external validity of laboratory findings for the context of interest, and complement alternative empirical approaches, including RCTs (for a taxonomy of field experiments, see Harrison and List, 2004).

1.4 Outline of the Thesis

To address these gaps, I apply three distinct research methods: (i) a theoretical analysis to study the interaction between reference-dependence and liquidity constraints; (ii) panel techniques to control for unobserved heterogeneity that confounds estimates of adverse selection; and (iii) a framed field experiment to unravel free-riding as a mechanism driving low take-up. The remainder of this chapter describes the three essays in Chapters 2 to 4. Chapter 5 summarizes the main findings, policy implications, and provides directions for future research. Chapters 6 and 7 conclude with a non-technical summary in English and Dutch, respectively.

Shattered plans? Save more. On reference-dependence and liquidity constraints.

Chapter 2 combines the fairly separated lines of literature on financial market imperfections and prospect theory, and analyzes how the two combined affect consumption plans. A large body

of literature suggests that behavior is better described by reference-dependent preferences as in Figure 1.3 than by standard risk aversion as in Figure 1.1. The evidence is however based mostly on experimental data (Wakker, 2003; Abdellaoui et al., 2007; Booij and Van de Kuilen, 2009). A few studies use empirical data from countries with high bank penetration, from New York taxi drivers' labor supply (Crawford and Meng, 2011) and sellers' behavior in US housing markets (Genesove and Mayer, 2001) to behavior in a Dutch game show (Post et al., 2008). Further, no existing theory predicts reference-dependent behavior in a context of imperfect financial markets, challenging the interpretation of tests for prospect theory outside the laboratory.

The chapter contributes by introducing financial market imperfections, modeled as a liquidity constraint, in Kőszegi and Rabin (2006, 2009)'s application of prospect theory with endogenous reference points. In this framework, consumers derive utility not only from consumption itself, but also from changes relative to initially planned consumption levels. The gain-loss utility from changing a plan is described by Figure 1.3. A consumer's initial plan now serves as the reference point r . By loss aversion, consuming less than planned hurts more than a gain is pleasant, and by diminishing sensitivity, it is optimal to spread gains over time, but concentrate losses in one period and carry through the plan for the other period. Finally, the model assumes that changes in plans regarding immediate consumption are felt more than changes in plans regarding future consumption.

The study analyzes whether this framework yields qualitatively different consumption plans than reference-independent frameworks in which consumers derive utility only from consumption itself. Another question is to what extent the results depend on the presence of a liquidity constraint. A key insight from the study is that when consumers are unable to borrow, reference-dependence might induce them to save *more*, while it would induce them to save *less* in a context of perfect financial markets. Without liquidity constraints, the reference-independent optimum is to smooth consumption. Reference-dependent consumers consume more in early periods because they attach higher value to changes in plans regarding immediate consumption. In other words, without liquidity constraints, reference-dependence reduces the optimal level of savings.

By contrast, with liquidity constraints, being unable to borrow might force consumers to give up present plans even when future plans are still attainable. Consider a consumer who plans to save because her income in the early period is relatively high. Just prior to consumption, she finds out that this income is lower than she thought it would be and insufficient to consume the level planned for the early period. A reference-independent consumer would decide to spread the loss and dissave in this case. But a reference-dependent consumer may decide to take the full loss immediately and keep saving in order not to give up future plans. Separately, neither a liquidity constraint nor reference-dependent preferences can account for this phenomenon.

This also suggests that reference-dependence by itself can explain low take-up of health

insurance as a protection against future health risk, but does not provide a convincing explanation in the presence of liquidity constraints. Conceptually, health insurance is a type of savings - it implies giving up consumption in early periods, with potential benefits accruing to the consumer in the future. Since reference-dependent consumers facing liquidity constraints are willing to save more than reference-independent consumers, they are also willing to pay more for health insurance. Thus, prospect theory does not necessarily explain low take-up in a context of imperfect financial markets. This novel insight, that is qualitatively different from predictions in the literature on prospect theory, highlights a strong interaction between behavioral decision-making theories and imperfect financial markets.

Intra-household allocations of micro health insurance

An alternative explanation for low take-up is adverse selection, meaning that potential beneficiaries only enroll when they expect to use more health care. Chapter 3 therefore tests for adverse selection within households, i.e. whether households are more likely to enroll members for whom they expect higher health expenditures. This study analyzes a micro health insurance program in rural Nigeria, where households have the option to enroll individual members. A major limitation of individual-based enrollment is that households may only enroll household members with high health risks, harming the sustainability of the risk pool.

The prior literature on adverse selection is inconclusive, but mostly analyzes programs with family-based enrollment. Hence, the existing literature does not look at insurance allocations within families and this chapter is the first to use intra-household variation in insurance status. With individual data for all household members, the data can be interpreted as a panel. The study exploits this structure, using household fixed effects to eliminate confounding unobserved heterogeneity at the household level. The aim is to shed light on whether households themselves consider health risk when purchasing insurance, by investigating whether differences in health risk influence the intra-household insurance allocation.

The study shows that many households purchasing insurance enroll a few household members only, especially when the household is large and relatively poor. But despite the opportunity to enroll only high-risk members, there is no evidence of adverse selection. In a control sample, where health insurance is unavailable and the price effect of having insurance does not bias health expenditures, baseline health is a significant risk factor for follow-up health expenditures. However, baseline health does not affect the intra-household insurance allocation.

Adverse selection is limited when conditioning on age and gender - risk factors that the insurance provider can price. Some adverse selection occurs in a program area with a public facility, but not in a program area where a community-based health care provider was actively engaged in enrollment. I show that adverse selection is limited both in terms of objective health indicators, as well as subjective health, in contrast to our priors. This suggests that adverse

selection cannot explain low take-up. In fact, it seems worthwhile experimenting with individual-based enrollment to enhance take-up among households unwilling or unable to pay the bulky family premium in one go.

This chapter uses data collected by the Amsterdam Institute for International Development (AIID) and Amsterdam Institute for Global Health and Development (AIGHD), commissioned by the Health Insurance Fund (HIF). I was involved in the impact evaluation of this scheme first as research assistant and later as AIID research coordinator. I assisted in the research design, developed survey instruments and provided technical support during the data collection. I coordinated both waves of data collection in Nigeria in 2009 and 2011 and liaised with the donor agency and implementing partners. Chapter 3 was inspired by experiences during the fieldwork, but not formally part of this appointment. The chapter does not necessarily reflect the views of the Health Insurance Fund or her implementing partners.

The social dilemma of microinsurance

A final hypothesis is that take-up remains low because the poor partly rely on assistance from their social network to cope with health shocks. Chapter 4, which is my job market paper, attributes low demand for health insurance in microcredit groups to a social dilemma. Microcredit clients typically share liability for group members' loans and contribute if necessary, for instance when a group member is ill and cannot repay her loan. But if too many group members are ill, the group will default and can no longer borrow. Insurance reduces the group default risk and may hence improve welfare. Chapter 4 analyzes insurance decisions within this institutional context.

The study uses non-cooperative game theory to show why jointly liable clients will forgo welfare-improving insurance. In their context, insurance is a public good. Buying insurance reduces the need for peers to contribute, but does not increase a client's private earnings. As a result, when clients can enroll in insurance individually, less risk-averse clients are tempted to free-ride on contributions from their peers and opt out of insurance. More risk-averse clients, willing to pay more for insurance, enroll as long as a sufficient number of peers do the same. Group insurance in which a group can enroll if and only if all group members are willing to join binds both types to the group optimum and eliminates free-riding.

The study tests this theory by means of a framed field experiment, closely mimicking the insurance decision in jointly liable microcredit groups. This field laboratory experiment with 355 microcredit clients in Dar es Salaam, Tanzania, incentivized the insurance decision in a controlled setting where distortions of e.g. initial beliefs, social ties and health do not bias the results. The experiment also provides insights into the dynamics of repeated insurance decisions within a short time span. Further, the framing enhances external validity compared to conventional public good games played by university students.

Participants played two games in the experiment. An initial individual insurance game

without joint liability elicits a measure of clients' risk aversion, labeling every participant as either more risk averse or less risk averse. Participants are then randomly assigned to a jointly liable microcredit group. Depending on the treatment, taking insurance is an individual decision or a group decision. In the individual insurance treatment, it is possible to free-ride and we predict low demand only among the less risk averse client type. Both types will enroll in the group insurance treatment if full group enrollment optimizes their welfare. This treatment serves as a benchmark to reveal jointly liable clients' preferences over insurance.

The experiment provides evidence of substantial free-riding. In the group insurance treatment, nearly all participants are willing to join health insurance. But in the individual insurance treatment, not even half of all less risk averse participants enroll. They free-ride at the expense of their more risk averse insured group members, who keep purchasing insurance and earn substantially less in the individual insurance treatment. When less risk averse clients are grouped with peers of the same type, individual insurance reduces demand even further relative to group insurance. This is consistent with our theoretical predictions. In the experiment, individual insurance is subject to free-riding. Group insurance eliminates free-riding, can reduce inequality within groups and potentially increases enrollment in health insurance.

The principal investigator of this framed field experiment is Wendy Janssens. I was involved in this project immediately after funding was secured and raised additional funding, allowing us to run additional treatment arms for future research on the demand for group insurance under joint versus individual liability. Based on discussions with Wendy Janssens, I developed the theoretical framework tested in the experiment. From February to April 2010, I coordinated the fieldwork in Dar es Salaam. We jointly conducted the analyses that led to a first draft paper. The chapter in this thesis is based on subsequent empirical analyses that I conducted mainly during a semester as visiting research student at the University of California Berkeley, where I drafted a new version of the paper. This draft, then, went through several rounds of revisions by Wendy Janssens and me.

In sum, reference-dependent preferences and adverse selection within households do not explain why take-up of micro health insurance remains low. When consumers face liquidity constraints, they are not always able to carry through their present plans, and reference-dependence can induce them to buy more insurance in order not to give up future plans. The case study in rural Nigeria finds that health risk does not determine intra-household insurance allocations. A more convincing explanation is the existence of risk-sharing arrangements, as shown by a framed field experiment with microfinance clients. This laboratory experiment in Tanzania provides substantial evidence of the theory that less risk-averse clients opt out of individual insurance because ill clients can free-ride on contributions from their microcredit group. Offering group insurance in which either all or no group member enrolls has the potential to solve low take-up.