UNITED NATIONS
UNIVERSITY

UNU-WIDER
World Institute for Development
Economics Research

Research Paper No. 2006/24

Applying Behavioural Economics to International Development Policy

C. Leigh Anderson[1]
and Kostas Stamoulis[2]

March 2006

Abstract

Many development policies and programmes are premised on a traditional economic model of rationality to predict how individuals will respond to changes in incentives. Despite the emphasis of these programmes on poverty reduction, economists and the development community in general are still unable to fully understand how the poor make decisions, especially under uncertainty and over time. Individuals avail themselves less than predicted in health programmes, participate less than expected in market opportunities, under or over insure themselves, and make short-run decisions that are inconsistent with their long-run welfare. The rise and fall of different descriptive models and paradigms of poor household behaviour can partly be attributed to this limited understanding. More helpful answers may lay within behavioural economics, that these insights are particularly important for poor populations, and that they can improve the future design, implementation and subsequent effectiveness of development programmes. Behavioural economics is an approach that rigorously …/

[1]University of Washington; [2]FAO.

combines the insights of psychology and economics to try to better understand and predict human decision making. Empirical evidence is helping us learn, for example, how cognitive limitations, fairness, loss aversion, framing of choices, variable discount rates, and the qualitative dimensions of risk—such as proximity and control—affect decision making. The regularity of many of these anomalies suggests that these behaviours are anomalous only to traditional models, but that they may otherwise be the norm.

Acknowledgements

The World Institute for Development Economics Research (WIDER) was established by the United Nations University (UNU) as its first research and training centre and started work in Helsinki, Finland in 1985. The Institute undertakes applied research and policy analysis on structural changes affecting the developing and transitional economies, provides a forum for the advocacy of policies leading to robust, equitable and environmentally sustainable growth, and promotes capacity strengthening and training in the field of economic and social policy making. Work is carried out by staff researchers and visiting scholars in Helsinki and through networks of collaborating scholars and institutions around the world.
www.wider.unu.edu   publications@wider.unu.edu

Introduction

Many development policies and programmes are premised on a traditional economic model of rationality to predict how individuals will respond to changes in incentives. Despite the emphasis of these programmes on poverty reduction, economists and the development community in general are still unable to fully understand how the poor make decisions, especially under uncertainty and over time. We do not understand why, for example, individuals do not regularly adopt subsidized technologies such as ventilated cooking stoves to reduce health risks. Even when cultural norms and other constraints are considered, our models do not explain why a massive HIV/AIDS information campaign does not encourage individuals to regularly use condoms. Individuals avail themselves less than predicted in health programmes, participate less than expected in market opportunities, under or over insure themselves, and make short-run decisions that are inconsistent with their long-run welfare. The rise and fall of different descriptive models and paradigms of poor household behaviour can partly be attributed to this limited understanding. We believe that some more helpful answers may lie within behavioural economics, that these insights are particularly important for poor populations, and that they can improve the future design, implementation and subsequent effectiveness of development programmes.

Behavioural economics is an approach that rigorously combines the insights of psychology and economics to try to better understand and predict human decision making. Empirical evidence is helping us learn, for example, how cognitive limitations, fairness, loss aversion, framing of choices, variable discount rates, and the qualitative dimensions of risk—such as proximity and control—affect decision making. The regularity of many of these anomalies suggests that these behaviours are anomalous only to our traditional models, but that they may otherwise be the norm.

The accumulated evidence is largely from the USA and Europe, with little comparable work from developing or transition countries. Most of the work has been laboratory based, with far fewer observations from the field. The motivation for trying to understand the preponderance and nature of similar behavioural anomalies in less developed countries stems from our prior that they are at least as prevalent, and that they will more acutely affect policy outcomes because there are fewer formal institutions to temper their effects. This prior stems from three developing country characteristics: the greater incidence of poverty and food insecurity, larger rural populations, and the dearth of well functioning markets.

We begin by discussing what we mean by behavioural anomalies and with a brief review of some USA and Western European experiments. We draw heavily on earlier experimental summaries, and in particular, Kahneman (2003), Thaler (1991), and Rabin (1996). In section two, using original field data collected with stated and revealed

preference surveys in Vietnam and Russia, we examine discount rate patterns, risk attitudes and decision heuristics. Our final section briefly summarizes why we believe these ideas are particularly important for international development, and what some of the obstacles are to change.

1   What are behavioural anomalies and behavioural economics?

Over the past twenty-five years economists have been increasingly focusing on the behaviours of individuals that deviate from what would be predicted from our standard model of rational maximization. The focus is not on behaviours driven by social and other institutional factors that have been neglected (or not properly captured) in economic models. Of interest are psychological factors that although susceptible to, or shaped by, the institutional environment, differ for other reasons among individuals living in identical circumstances.

Behavioural economics is an effort to understand how systematic these deviations are and to adjust our models accordingly. Our standard economic model posits choice as being the outcome of an individual maximizing stable, well behaved preferences, U(x), over a set of goods and services, subject to a set of measurable constraints such as income or prices. In this paper we examine four potential model misspecifications that we believe are particularly problematic if they arise in poverty alleviation policies and programmes. The first is the common discounted utility model (DU), because of how fundamental intertemporal choices are to growth, borrowing, investment and all sustainability issues where costs and benefits are spread over time. The second is the dominant expected utility model (EU), because the poor, particularly those who live in rural areas and rely on agriculture, face considerable risks and generally have fewer formal institutions to manage those risks. The third is the disregard by the prevailing economic model for reference levels and fairness, because any intervention changes the status quo. Our forth concern is with the method behind the model, and the assumption that 'Every decision is thoroughly contemplated, perfectly calculated and easily executed' (Mullainathan forthcoming). Instead we look at decision heuristics, because information to inform decisions is particularly complex in many developing and transition economies and scarce in rural environments—see Kreps (1990) and Rubinstein (1998), among others, for a technical exposition on which traditional behavioural axioms are violated by observed anomalies.

1.1 Decision making over time

Discount rates affect investments in education and health, use of the environment, borrowing and saving, and all choices where something has to be given up in the present for a return in the future. Since Paul Samuelson's seminal article in 1937, intertemporal choices have been most commonly modelled with discount rates that are invariant to the period in which the choice occurs. The assumption is that individuals have constant discount rates, that is, they will make the same trade-offs within any period regardless

of the proximity of that trade-off. This suggests that the short-run discount rate we experience in postponing immediate consumption is the same as the long-run discount rate we use in planning future trade-offs (Harris and Laibson 2001).

Yet, as Rabin (2002: 669) writes, 'Common sense, millennia of folk wisdom, and hundreds of psychological experiments all support present-biased preferences'. For most of us, the cost of foregoing something today for tomorrow is higher than the cost of agreeing to give up something ten days from now for a return on the eleventh day. Many economists believe that intertemporal choices are better represented by a hyperbolic or quasi-hyperbolic discount function than the more commonly employed exponential function (Rabin 2002; Rabin 1996; Laibson 1997; Loewenstein and Elster 1992; Loewenstein and Prelec 1992; Ainslie 1991). The evidence suggests that discount rates do vary with time, and that short-run discount rates, experienced at the moment, are higher than the long-run discount rates we project forward into our planning horizon. Hence people with time varying discount rates, if they lack perfect self-control, may pursue short-run actions that they had previously calculated were not in their best long-run interest—they may consume their savings, fail to stick to a debt repayment schedule, skip school or healthcare visits, or choose environmentally unsustainable production methods.

1.2 Decision making under uncertainty

The traditional model of individual choice under uncertainty represents preferences by an expected utility function, the shortcomings of which have been well documented (Allais 1953; Kahneman et al. 1982; Machina 1987). Experiments instead support prospect theory, with a value function defined over changes in one's position rather than wealth, and which is assumed concave over gains, convex over losses and more steeply sloped over losses than gains (Kahneman and Tversky 1979). These features better account for the impact of losses and gains and loss aversion—individuals being risk averse over gains but risk seeking over losses. Unlike basic utility theory, research suggests that losses hurt more than commensurate sized gains help—in several cases with losses being weighed more than twice as heavily as gains (Knetsch 1995). Prospect theory seems especially relevant for developing countries and rural economies where income levels can be highly variable.

The subjective weighting function of prospect theory reflects three characteristics of observations on decision making: the overweighting of small probabilities and underweighting of large ones; decreasing relative sensitivity that discounts probabilities further from one proportionately less than probabilities closer to one (sub proportionality); and increasing absolute sensitivity towards the endpoints of probabilities equal to zero and one (subadditivity) (Prelec 2001: 86).

In addition to magnitude, the qualitative dimensions of risky outcomes have been shown to affect risk perceptions. The characteristics that shape decision making about risk

include magnitude, reversibility, control, familiarity, proximity and distribution of impacts, among others (Slovic et al. 1979; Pate 1983; NRC 1996). The risk perception literature shows that individuals systematically overestimate the size of risks that are small (the Allais paradox), unfamiliar, involuntary and uncertain (the Ellsberg paradox). By contrast individuals underestimate the size of risks that are more certain, larger, familiar, or in some sense voluntary.

1.3 The true form of U(x): reference levels and fairness

Most of us understand that behaviour is motivated by considerations other than pure self-interest. What we understand less well is how, and when, to weight these motivations. If self-interest is narrowly considered as increasing material wealth, for example, how much wealth will one forego to satisfy fairness, revenge or altruism? And when are the motives governed by more complex issues such as reciprocity? For example, studies suggest that people are more likely to contribute to public goods when they believe others are also doing this, even though from an efficiency perspective there are diminishing social benefits to each additional person contributing. Pure altruism would suggest that one should give more when others give less (Rabin 1996:13).

Repeated experiments suggest that individuals are willing to suffer monetary losses to punish opponents for outcomes that they perceive as unfair (see for example, Rabin 1996, and Camerer and Thaler 1995) They may either judge the resulting distribution or the intent of the distribution as unfair (Bereby-Meyer and Niederle, 2005) In USA experiments, individuals will refuse allocations that deviate much from 50-50 if their refusal means that both players receive nothing. That is, they are willing to give up their smaller share to prevent the other player from receiving their larger one. Interestingly, in the results of a study of ultimatum games in fifteen small-scale societies there was a positive correlation between offers to share more equitably and degrees of market integration and co-operation (Henrich et al. forthcoming). Under what conditions and forms of economic organization, if any, can we regularly predict that satisfying the desire to retaliate will outweigh the desire to increase wealth?

Kahneman et al. (1986) conducted several experiments to assess how fairness might be a constraint on profit seeking in the market. Consider the different percentage of the approximately 125 respondents who considered a company decision to be unfair under the circumstances in Q1 and Q2:

> Q1. A company is making a small profit. It is located in a community experiencing a recession with substantial unemployment but no inflation. There are many workers anxious to work at the company. The company decides to decrease wages and salaries 7 per cent this year. Sixty-two per cent considered this unfair.

> Q2. A company is making a small profit. It is located in a community with substantial unemployment and inflation of 12 per cent. There are many workers

anxious to work at the company. The company decides to increase salaries only 5 per cent this year: Twenty-two per cent considered this unfair.

From their experiments they proposed the principle of reference transactions and dual entitlement: that buyers feel entitled to the terms of some reference transaction and firms to some reference profit, without necessarily knowing profit levels at the reference transaction. Transactors found it largely unacceptable for firms to raise prices and appropriate surplus from demand increases, but acceptable for firms to pass on cost increases. This included price or wage changes in response to commodity shortages or labour surpluses. If these views of fairness affect individuals' decisions on where, and with whom, to exchange, or constrain profit seeking by vendors in the market, then we should expect that people will respond differently than traditionally predicted to policies that regulate or liberalize markets.

Reference levels apply to wealth as well as prices. Considerable evidence suggests that individuals are sensitive not only to absolute wealth levels, but also to changes in their relative wealth. The baseline could be their own past wealth, such that it is the increase or decrease in their wealth that matters as much as the final level. Or it could be changes in the wealth of a neighbour, friend, or other individual they are apt to compare themselves to. Thus, for example, an individual could feel worse off from the good fortune of their neighbour or from a colleague getting a pay raise, despite no change in their absolute wealth levels.

Many of these ideas on relative wellbeing stem from the literature around happiness. Beginning with Easterlin (1976), several studies have shown that although on average individuals in wealthier countries are happier than their poorer country counterparts, average happiness levels do not rise as countries grow wealthier. One hypothesis is that as average wealth levels rise, individuals do not necessarily feel relatively better off. Graham (2004) explores some of these ideas in the context of globalization, growth, and inequality for developing countries. She cites results from Gurr (1970), who after studying conflict-related deaths in over a hundred countries, cites relative, not absolute, deprivation as 'the basic, instigating condition for participants in collective violence...' (Graham 2004: 5).

Clearly, the extent to which it is relative income that affects welfare must be of some concern to development economists. First, judgements of fairness may underlie individual responses to all mechanisms that allocate scarce resources. Depending on the principles underlying how programme resources are allocated—maximizing social welfare, helping the most needy, or helping everyone equally—judgements of fairness may affect programme participation. Ignoring perceptions of fairness may affect interventions designed to make markets work for the poor—including increasing access to and allocating resources through vouchers, subsidies and other market-based mechanisms. Second, the goal of economic growth presumes wellbeing rises with absolute income and is neutral with respect to inequality, but that rising inequality may

affect the pace of growth. This evidence suggests that the concern with inequality should rest directly with its effect on wellbeing, not just the long-run growth process. For the very poor, however, changes in absolute income may still be the most important for wellbeing. McBride (2001) estimates for a developed country sample that relative income effects are smaller at lower income levels.

Loss aversion and mental accounting

Loss aversion may contribute to two other phenomena called the status quo bias and the endowment effect. Experiments suggest that individuals tend to prefer the status quo to changes that involve losses in some dimension, even when these losses are coupled with gains in another direction. In other words, actions that are, on net, revenue neutral or even revenue enhancing may be welfare decreasing if they involve losses. And either because of the status quo bias or inertia, default options dominate choices. A study of insurance in Pennsylvania and New Jersey found that with full coverage as the default option in Pennsylvania the take-up rate was 79 per cent. With limited coverage as the default option in New Jersey, only 30 per cent of drivers choose full coverage (Johnson et al. 1993). We should expect project recipients to be willing to take bigger gambles to maintain the status quo than to acquire it in the first place, and that default options will dominate choice for reasons other than traditional welfare maximization.

The increased value attributable to possession or property rights, commonly called the endowment effect, was documented in a series of experiments with University of Victoria undergraduate students. In one experiment, respondents were divided into three classes. One class was given a choice between a mug of value $4.95 and a chocolate bar of approximate value $6.00. Students were not told the purchase price of the items. Student in the first class were given a choice between the two items and 56 per cent chose the mug over the chocolate, giving us a sense of a distribution of preferences. Students in a second class were first given the mug, and then five minutes later given the opportunity to trade for the chocolate—89 per cent chose to keep their mug. Students in the third class were initially given the chocolate and then the offer to trade for a mug. This time, 90 per cent chose to keep their chocolate. (Kahneman et al. 1990).

Evidence suggests that individuals have mental accounting systems that influence decision making in ways not predicted by traditional economic theory. Thaler has documented numerous examples, including how individuals are more likely to spend 'windfall' winnings on luxury items than equivalent sized salary increases. People often give gifts that recipients value but would not purchase for themselves. Many individuals seem to have mental accounts for different expenses, such as entertainment and education. Consider Thaler's example revealing the discrepancy in choices over buying a theatre ticket that have the same net impact on wealth:

Q1. Imagine that you have decided to see a play where admission is $10 per ticket. As you enter the theatre you discover that you have lost a $10 bill. Would you still pay $10 for a ticket to the play? Yes: 88 per cent. No: 12 per cent.

Q2. Imagine that you have decided to see a play and paid the admission price of $10 per ticket. As you enter the theatre you discover that you have lost your ticket. The seat was not marked and the ticket cannot be recovered. Would you pay $10 for another ticket? Yes: 46 per cent. No: 54 per cent.

The bundling and sequencing of monetary gains and losses also matter. With equivalent cash outcomes, people are happier winning two small prizes ($50, $25) than one large prize ($75); paying one large penalty (-$150) than two small penalties (-$100, -$50); paying one large penalty and winning one small one (-$200, $25) than paying one medium sized penalty (-$175); and winning one small prize ($20) than winning one large prize and paying one medium penalty ($100, -$80) (Thaler 1991: 55). Experiments in health indicate that with the same net outcome, people prefer improving sequences to declining sequences and that the duration of a sequence of events is less important than the experience of the final frame (Chapman 2003).

Economists have long argued that money is the best gift because it can be redeemed on whatever the recipient wants, and that the most efficient redistribution is via lump sum transfers to the poor. These examples suggest that how money is received, and what it is mentally ear marked for, matter both in terms of the recipients' welfare and their decision on how to use the resources. These findings have implications for debates over, for example, the effects of vouchers versus cash, earmarking funds, and segregating gains and bundling losses. Loss aversion, the endowment effect, and mental accounting may explain current challenges to how we conceptualize opportunity and sunk costs— people often behave as though opportunity costs matter less than out-of-pocket expenses.

1.4 Systematic errors in maximizing U(x): decision heuristics

The expected utility model is particularly inappropriate for complex and infrequent decisions, where evidence supports a bounded rationality model. Individuals often resort to decision heuristics that can produce behaviour that appears imperfectly rational. Decision heuristics refer to the simple rules of thumb that individuals use to make decisions either in the absence of full information or when they are unable or unwilling to process all the information that is available. Some common heuristics are representativeness, availability, anchoring, and affect.

When people use the representativeness heuristic, their judgement of the probability that one event or person originates from or belongs to another class is based on how the events or people resemble each other. The representativeness heuristic judges such a frequency by comparing the similarity of the case with the image or stereotype of the

class, often to the exclusion of prior probabilities, base-rate frequencies, sample size, regression to the mean, and other factors that should affect judgements of probability (Tversky and Kahneman 1974). For example, people tend to over estimate how often a small group will closely resemble the parent population. The bias occurs when frequency and similarity are not well correlated. As an example from Thaler (1991: 153, originally from Tversky and Kahneman 1983):

> Consider a regular six-sided die with four green faces and two red faces. The die will be rolled 20 times and the sequence of greens (G) and reds (R) will be recorded. You are asked to select one sequence, from a set of three, and you will win $25 if the sequence you choose appears on successive rolls of the die. Please check the sequence of greens and reds on which you prefer to bet:
> A. RGRRR
> B. GRGRRR
> C. GRRRRR

Since A is a subset of B, it must be more probable than B. But B may appear more representative of a probable sequence because of the two Gs. In this experiment, 63 per cent of respondents chose B, and 35 per cent chose A.

The availability heuristic leads individuals to overestimate probabilities of recent or vivid events. It arises when people estimate the frequency of a class by the ease of recalling specific instances in that class. People also use anchoring heuristics, where arbitrary amounts become the bases for forming numerical estimates of uncertain quantities. One example is an experiment asking subjects to state the percentage of African countries in the United Nations after being given an arbitrary starting point (based on a ball thrown on a spinning wheel with numbers from 1 to 100). The median estimates of membership were 25 per cent for groups with ten as a starting point, and 45 per cent for groups that received sixty-five as a starting point (Tversky and Kahneman 1974).

The affect heuristic, recently proposed by Slovic and others (2002), describes the bias estimating probabilities that results from an individual's like or dislike of an outcome. They argue that the affect heuristic can affect one's evaluation of costs and benefits and even the predicted economic performance of various industries (Kahneman 2003: 1463).

Finally, choices between alternatives can also be affected by the context of the event and the way in which a risk is framed (Kahneman and Tversky 1979). For example, McNeil et al. (1982) framed a hypothetical choice about lung cancer treatments from surgery in two ways; as either a 68 per cent chance of survival or a 32 per cent chance of not surviving. With a 68 per cent chance of survival, 44 per cent of respondents chose surgery over radiation. Framed as a 32 per cent chance of not surviving, the number dropped to 18 per cent.

2   Developing country studies

Our goal is to understand the prevalence and nature of anomalies in developing country contexts in the field (as opposed to the laboratory), and to compare decisions across policy domains. Examining behavioural anomalies first requires designing and implementing stated and revealed preference surveys for eliciting and measuring characteristics that affect individual decision making. Stated preference methods involve asking for responses (by choosing, ranking, or providing an open-ended answer) to sets of hypothetical scenarios, defined by underlying attributes. For example, to elicit a discount rate, a stated preference experiment might ask a respondent to choose between an immediate lump sum payment and a series of constant payments over a number of years (Cameron and Gerdes 2002) or to provide an amount of money that would equate lump sum payments to be received at different times (Anderson et al. 2004).

Revealed preference questions measure actual behaviours. For example, to study risk attitudes, a stated preference question may ask a respondent to indicate their preference for a different set of hypothetical gambles with a coin toss, the return that would make a certain outcome equivalent to a gamble, or how important they believe risk taking is for certain outcomes like financial success. Conversely, revealed preference surveys would look at actual choices that reveal risk tolerance, for example, at how the individual has diversified their income sources or other portfolios. Binswanger (1980) looked at the relationship between stated preference and actual behaviour.

Common concerns with stated preference surveys are validity and learning. Sceptics worry that especially without remuneration respondents will put little effort into responding. In a review of over 70 experiments, however, Camerer and Hogarth (1999) conclude that overall there is no effect on mean performance though response variance declines with financial incentives. Performance of some specific tasks, including complex ones, does improve with remuneration. The authors note (p.7), however, '…that no replicated study has made rationality violations disappear purely by raising incentives'. Though learning may occur in repeated experiments, experimenters have found that amateurs and experts in the field tend to perform similarly.

The vast majority of the methodological work on stated preference surveys has been in developed counties. Elements of these instruments can be adapted for use in developing countries, but experience suggests that great care is necessary at both the survey design and data analysis phase to control for the different context (Cameron and Gerdes 2002; Corso et al. 2001). For example, Kuechler (1998) has argued that a country's recognition of freedom of speech and culture of individualism will affect survey responses. Certainly in Russia and Vietnam one can imagine that many individuals still fear answering questions or offering opinions that might deviate from official expectations. In this section we briefly describe some ongoing work: discount rates in Vietnam and Russia; risk perceptions in Russia; and decision heuristics by level of

policymaker in Vietnam. Full details on these experiments can be found in Thaler (1981), Benzion et al. (1989), Anderson et al. (2004), and Anderson et al. (2005).

Example 1: intertemporal choices in Vietnam and Russia

There is a considerable literature on time preference and discount rates—the rate at which individuals trade off future for present consumption. One type of experiment to reveal discount rate patterns compares discount rates calculated from respondents comparing smaller, earlier rewards (or penalties) to larger, later rewards (or penalties). In 1981, Richard Thaler asked University of Oregon students to state the amount of money they would require to either postpone a fine or expedite the receipt of lottery winnings, assuming no risk. Thaler asked the question for lottery winnings of $15, $250, and $3000 and three month, one year, and three year delays. Several years later, Benzion et al. (1989) conducted a similar study with students at the University of Haifa and the Technion-Israel Institute of Technology. Their questionnaire asked for intertemporal choices over four scenarios of postponing or expediting a receipt or payment, time delays, and sizes of cash flow. The Israel study used dollar amounts of $40, $200, $1000, and $5000 and time delays of five months, one, two and four years.

For both the USA and Israel we show their results in Table 1 for the scenario that most closely corresponds to the Vietnam and Russia experiment. The first and fifth column represent the different cash flows in the USA and Israel, with the different time delays for repayment across the top row. The discount rates inferred from their dollar responses appear in the cells; median rates for the USA and mean rates for Israel.

Table 1: Inferred discount rates, USA (median) and Israel (mean)

Original amount, USA	Later amount paid in			Original amount, Israel	Later amount paid in			
	3 months	1 year	3 years		5 months	1 year	2 years	4 years
$15	277	139	63	$40	53.5	33.0	26.5	20.6
$250	73	34	23	$200	32.1	23.6	21.0	15.7
$3000	62	29	23	$1000	31.0	21.9	16.6	16.3
				$5000	26.1	19.2	14.9	13.6
N = 20				N = 204				

Source: Thaler (1981), and Benzion et al. (1989).

In both the USA and Israel, the results suggest that discount rates are not constant, and that they decrease as the size of the cash flow increases. That is, the smaller the postponed fine or delayed receipt, the higher the discount rate. Discount rates also vary inversely with the period of time until repayment or receipt. The discount rate of 277 for $15 in three months is the highest in the US, and the discount rate of 53.5 for $40 in five months is the highest in Israel. The Vietnam and Russia studies were undertaken more than ten years later. In 2000, interviewers from Vietnam's Institute of Sociology randomly sampled individuals from two communes near Hanoi city: the urban Quynh

Mai and the more rural Thach Ban. In 2002 we replicated the Vietnam study with a random sample from in and around two southern Siberian cities.

We sought to replicate the USA and Israel questions as closely as possible, but we encountered several difficulties during pre-tests. First, respondents had strong feelings about receipts or payments from public institutions—often either extreme distrust or extreme allegiance. Non-governmental organizations (NGOs) elicited less emotion. Second, respondents had some difficulty with the idea of hypothetical trade-offs. Third, the discount rate for respondents in Vietnam and Russia fell to almost zero after three months; respondents were unable or unwilling to differentiate among longer time periods. Finally, we could not use comparable monetary amounts, even adjusted for purchasing power parity. Respondents in Vietnam were unable or unwilling to differentiate among large amounts beyond about half their annual income.

Ultimately in Vietnam we asked respondents to imagine that they had the opportunity to receive a loan from a NGO and that they had the choice of paying back the loan immediately or postponing the payment to a later date, at which time they would have to pay back a larger amount. The questions used sums of Vietnamese Dong (VND) 100,000, 1,000,000, and 4,000,000, and time periods of one day, three months, and one year. At the time of the study, US$1 was worth about VND 14,500, so the survey amounts were worth about $7, $70, and $276.

Russian respondents were asked to imagine that they had just received a loan with the choice of paying it back immediately or postponing repayment to a later date, at which time they would have to pay back a larger amount. We used sums of Russian roubles (RUR) 1500, 6000, and 30000 for the same Vietnam time periods of 1 day, 3 months and 1 year. At the time of the survey, US$1 was worth about RUR 30, so survey amounts were worth about US$50, $200, and $1000.

Table 2 reports the results the results for Vietnam and Russia. We report median rates for the Russian (and USA) study, where the size of the inferred discount rates and the standard deviation of responses was much larger, and mean rates for Vietnam (and Israel). Thaler speculates that this large deviation may be due to the hypothetical nature of the study or the age of the respondents. There are, however, cases of similar results in studies without hypothetical questions Hausman (1979), and Benzion et al. (1989) compute considerably lower rates with a similarly young sample. Nonetheless, as Thaler points out, what matters is the relative, not the absolute levels.

Table 2 again reports the varying dollar amounts for the two studies in columns one and five, the time periods across the top rows, and the discount rates in the cells within. Our results suggest that despite field and experimental differences and despite the vastly different demographics and circumstances of the respondents, the discount rate patterns are all inconsistent with assumptions of time invariant discount rates. The variance within the discount rates in each study all suggest the same conclusion: respondents'

discount rates are not constant over time, but rather vary inversely with time and the size of the cash flow, and in contradiction to the standard DU model, which would predict the same rate for one day, one month, and one year.

Table 2: Inferred discount rates, Vietnam (mean) and Russia (median)

Original amount, Vietnam	Later amount paid in			Original amount, Russia	Later amount paid in		
	1 day	1 month	1 year		1 day	1 month	1 year
$7	66.9	2.5	1.5	$50	102.4	3.2	1.3
$70	33.5	0.9	0.7	$200	56.2	2.4	1.1
$276	18.3	0.7	0.6	$1000	55.8	1.8	0.9
N = 232				N = 417			

Source: Anderson and Cullen (2005).

Our surveys also collected some basic information on respondents. Living rurally was a strong indicator of higher discount rates, which may partially reflect risk. Consistent with other studies, we found discount rates decreased with age and were negatively associated with income (Davies and Lea 1995; Anderson and Nevitte forthcoming). Irving Fisher (1930: 73) asserted that 'a small income, other things being equal, tends to produce a high rate of impatience'. Lawrence (1991: 54) found similar evidence of this in the United States. Her results may imply that impatience leads to poverty, as individuals with high rates of time preference choose jobs with low and flat pay scales rather than ones that pay well only after a period of training or education. Alternatively, poverty breeds impatience from living at or near subsistence. Relative, rather than absolute, income may also matter in ways we have not picked up. Even within a poor, rural commune in Vietnam, for example, members will distinguish between who is 'poor' and who is not.

Example 2: risk perceptions in Russia

Our Russian data were collected during a three-month period in the summer of 2000. Russian-speaking USA graduate students teamed with Russian graduate students to survey approximately 500 residents of both Novosibirsk and Irkutsk oblast in south-western Siberia.

Novosibirsk grew to prominence during the Second World War when, for security reasons, the Kremlin decided to relocate its military-industrial complex and the Academy of Sciences there. With over two million inhabitants, Novosibirsk is Russia's third largest city after Moscow and St Petersburg. The city hosts the arts, industry, thirteen institutes of higher education including its 'Academgorodok' university campus, and the seat of regional government (Carver 2003: 9). Irkutsk lies further to the east, near Lake Baikal, and has a history dating from the mid 1600s as a trading juncture between China and Russia's south. In the mid 1800s, the exile of several intellectual radicals to Irkutsk associated Siberia with labour camps and prisons. In 1898, the trans-

Siberian railroad added to its importance as a trading post. It is populated with small businesses, traders, and its 650,000 inhabitants are more ethnically mixed than Novosibirsk.

The survey contained multiple stated preference questions that were intended to measure risk and related psychological parameters: three coin toss questions, nine discounting questions, two risk and optimism attitudinal questions, and six questions on uncertain outcomes intended to represent qualitative dimensions of risk. We report only a few preliminary results here.

The coin toss questions took the usual form, with even odds for flipping heads or tails. Situation 1 and 2 involved simple prospects: a sure positive outcome for option one versus gambling for a positive pay-off in option two. Situation 3 was a binary prospect, offering two gambles, both with a possible loss. Pay-offs in roubles, the percentage of responses for option one (%) and total sample size (N) are in Table 3. At the time, US\$1 = 30 roubles.

Table 3: The expected values of the options in the three scenarios

	Option 1	%	Option 2	N
Situation 1	10,000	62	EV = 0.5(20,000)+0.5(0) = 20,000	948
Situation 2	300	52	EV = 0.5(450)+0.5(0) = 225	914
Situation 3	EV = 0.5(900)-0.5(30) = 435	76	EV = 0.5(6000)-0.5(4500) = 750	890

Source: Anderson and Cullen (2005).

Sixty-two per cent of respondents chose the riskless option in a fair gamble (option 1 in situation 1). More people were willing to gamble when the expected value of the gamble was less than the sure bet, as in situation 2. But in situation 2 the amounts are small, even in Siberia 300 roubles (\$10) is not a lot of money. Hence we expect that the small magnitude of the gamble is affecting this result, that is, people are more willing to give up a sure 300 roubles to gamble with smaller amounts (Rabin 1998). The results in situation 3 are consistent with loss aversion. People are willing to gamble for a much lower expected value in order to avoid a potentially large loss.

The question on the qualitative dimensions of uncertain outcomes was primarily intended to assess if risk perceptions varied by experience, that is, if small business owners perceived certain risks differently than non-small business owners. Arguably some of these outcomes, such as mafia involvement, taxes, and debt, represent other dimensions that have been found to affect risk perceptions, such as proximity, control, dread, and familiarity. The survey asked respondents to indicate which, of a number of listed reasons, they thought might lead someone to not be interested in borrowing. Responses were coded from a zero for strongly agree, to a three for strongly disagree. The OLS results (equally weak but easier to interpret than an ordered probit estimation) appear in Table 4. The dependent variable is the degree to which respondents disagreed

with proposed reasons for individuals not wanting to borrow money to start a business, hence the higher the response of the individual the more likely they disagreed. Income was excluded because of our concerns with its validity and we acknowledge the possibility that small business ownership is endogenous, biasing the estimates.

Table 4: Perceptions of the qualitative dimensions of risky outcomes

Dependent Variable	Taxes too high	Mafia is a problem	Worried about debt	No trust in legal system	Interest rate too high
Constant	0.553***	0.989***	0.679***	1.28***	0.985***
Age	9.74E-05	-0.007**	0.001	-0.011***	-0.006**
Gender, male = 1	0.314***	0.103	0.059	0.093	0.105
Education	0.042*	0.041	0.030	-0.016	0.009
City of interview Irkutsk = 1	-0.053	0.272**	-0.149*	-0.032	0.089
Internet access yes = 1 no = 0	0.027	0.122	0.155*	0.010	0.037
Small business owner, yes = 1	-0.174**	0.253**	0.004	-0.020	-0.151*
N = 610	$R^2 = 0.03$	$R^2 = 0.05$	$R^2 = 0.01$	$R^2 = 0.03$	$R^2 = 0.03$

Note: ***significant at 1%, **significant at 5%, *significant at 10%.

Source: Anderson and Cullen (2005).

In general, age is a predictor of risk attitudes: the older one is, the more likely they are to believe that concerns about the mafia, legal system, or interest rate would deter someone from small business ownership. Mafia is more of a concern in Irkutsk, debt is less (despite earlier results that discount rates are significantly higher). Gender is only significant in the case of concern over taxes, and internet access in the case of the mafia. Small business owners see these risks differently from the population as a whole. They are more likely to find taxes and interest rates too high, and less likely to view mafia involvement as a deterrent. These results are preliminary, but suggest that loss aversion matters and that familiarity affects subjective risk perceptions.

Example 3: fairness and decision heuristics in Vietnam

To understand systematic and recurrent biases in resource allocation and programme design requires taking these experiments a step further. It requires identifying regular differences in patterns of bounded rationality along the policy chain: between the policymakers, programme/project designers, and programme/project recipients. Policymakers convey broad ideas, directions and priorities, and thereby a flow of resources to intermediary groups—government agencies, quasi or non-governmental organizations, or members of the private sector—for programme design and implementation. These intermediaries convert policy statements into a set of rules that represent the incentives—constraints and opportunities—faced by recipients. Hence

14

decisions at each level allocate resources and frame the decisions for the group that follow.

Kahneman (2003) describes two modes of thinking: intuition and reasoning. Effort and association are two characteristics that distinguish these processes: intuition is effortless and associative, while reasoning is slow, effortful and rule-governed. We posit that the amount of effort one chooses to expend making a decision is a function of the responsibility they bear and experience they have over decision making and the outcome, and the expected net value of the outcome. The expected net value of the outcome depends on the decision makers' risk attitudes, and their expectations about the benefits relative to the effort, time, and other resource costs of making the decision. These costs are affected by their access to information, both internal (through association) and external (through exposure). Benefits depend on the probability of the outcome occurring and the qualitative dimensions of risky outcomes: how much control they perceive over it, their familiarity with it, and its regularity, proximity, and reversibility. If this theory holds, then the degree to which one expends effort and the type of decision heuristics they use is expected to vary by experience, responsibility, and risk perceptions. For most developing and transition countries, these characteristics can be expected to differ significantly between the policymakers and programme recipients. These parameters are also assumed to depend on the decision making domain. For example, one may be very conservative in financial decisions, but reveal a high risk tolerance in health decisions or a strong willingness to experiment with a variety of seeds. These domain sensitivities may split according to decision making roles, which are often segregated by age and gender.

To explore these ideas in Vietnam, we asked the same set of behavioural questions to forty relatively poor farmers (mean monthly income = 842,062 VND, or approximately $53) and to forty middle- and high-level professionals involved in policy making (mean monthly income = 4,210,213 VND, or approximately $265). We refer to this latter group as policymakers, though their direct involvement varies. Respondents were asked a series of questions designed to reflect cognitive effort and views on fairness. We report on some preliminary results. First, farmers and policymakers would allocate development funds in quite different proportions, with farmers allocating much more to healthcare and policymakers to the financial sector. Both favoured vouchers as a form of aid, but farmers were more in favour of unconditional money transfers than policymakers were. Policymakers would allocate more money to subsidized seed, while farmers would allocate relatively more to subsidized medicine and subsidized credit.

Following on the work of Kahneman et al. (1986), respondents were asked how fair they considered interest rate changes in response to an increase in demand versus an increase in costs of the lenders. On average, both groups considered it more fair for savings and co-operatives to raise rates than for moneylenders to do so. This distinction, perhaps driven more by views about the lenders than by behavioural issues, is more

pronounced that differences according to whether the increases are demand versus cost driven. One interesting finding is that policymakers were significantly more inclined to give an answer and state whether they thought the increases were acceptable or unfair. Almost half of the farmer sample responded that they did not know, compared to approximately 10 per cent of policymakers.

Respondents were given a series of questions first suggested by Frederick et al. (2002, reported in Kahneman 2003) such as: medicine and a vaccination cost VND 110,000 in total. The medicine costs 100,000 more than the vaccination. How much does the vaccination cost? In Frederick's experiments, 50 per cent or fewer university students gave a correct response to this question. In our sample, the per cent of correct and incorrect responses are given in Table 5. The questions differed slightly in their wording, though the algorithm for determining the answer is the same in all cases.

Table 5: Per cent answering correctly

Decision	Socks and shoes	Medicines and vaccines	Chicken and rice	Interest and loan fees
% correct farmers	5	0	25	0
% correct policymakers	70	60	90	45

Source: Anderson et al. (2006).

Table 5 shows that for each decision a much higher per cent of policymakers answered correctly. For example, 70 per cent of policymakers compared to 5 per cent of farmers in the question that substituted socks and shoes for medicine and a vaccine. The questions appeared in the survey in the same order as in the table. Given the common algorithm, the variation by domain is surprising. In all cases, policymakers also took at least a minute longer (significant at 10 per cent) to answer.

We then asked respondents to make choices among outcomes with multiple attributes: a medicine for diarrhoea, a modern variety seed, and a loan officer for a local savings co-operative. In the case of the seed, for example, they were given the rankings out of 10 for three varieties:

> Variety 1 – 4, 5, 9
> Variety 2 – 6, 6, 6
> Variety 3 – 10, 10, 0

The 4, 5, and 9 refer to the ranking out of 10 for each of three desirable attributes (taste, pest resistance and drought resistance), but respondents were told that it was unknown which ranking went with which attribute. In all domains the farmers were more likely than policymakers to choose the option where a single element and a subset of attributes scored the highest, but where one attribute scored 0 (variety 3 in this case). We interpreted the option with a zero as the riskiest, but in two batteries of risk questions— one with lotteries and one asking to identify their own risk preferences—farmers were

significantly more risk adverse than the policymakers. This suggests that something other than risk preferences may be driving the decision heuristic on option choice. Policymakers were significantly more likely to choose the option where every attribute scored equally (variety 2 in this case), and in all cases but one it was their first choice. The one exception was in the seed question where their first choice was the last choice of farmers. This is arguably the domain where direct experience differs most between policymakers and recipients.

In many cases the differences emerging between farmers and policymakers is striking. It is premature to know whether and how these results derive from differences in cognitive processes, and if so, what is behind this—exposure, experience, a sense of control, responsibility, etc. But it is not, we believe, premature to recognise that if systematic differences are revealed in these behavioural experiments that they may lead to a disconnect between the resource allocations and policy designs of policymakers and programme recipients.

3   Why does this matter for international development policy?

Three developing country characteristics suggest that behavioural anomalies may be even less anomalous (i.e. they may be much closer to the norm) than observed in the USA and Europe. These characteristics are:
1.   greater incidence of poverty and food insecurity;
2.   lower incidence of well functioning markets, combined with greater price and output variability common to rural agricultural markets and aggravated by the occurrence of extreme events such as natural disasters and war; and
3.   greater rural proportion of the population and the concentration of poverty in rural areas.

The poor may disproportionately represent populations for whom traditional behavioural assertions are inappropriate if anomalies arise from threshold effects around base levels such as minimum subsistence income. And more individuals in developing countries live in poverty, with few assets, and are food insecure. The 2004 gross national income per capita was US$510 for low-income countries compared to US$32,040 for high-income countries (World Bank 2005; the gap narrows using figures adjusted for purchasing power parity (PPP) to US$30,980 and US$2,290).

Critics of behavioural economics argue that the anomalies of imperfect rationality can be ignored because the actions of a few that deviate from utility or profit maximization will be eliminated by arbitrage and competition in well functioning markets. But much of the exchange that occurs in developing and transition economies takes place in small, personal, informal, and poorly functioning markets, such that any anomalous behaviour is more likely to influence resource allocation. Schneider and Klinglmair (2004) estimate the average size of shadow economies as a percentage of official GDP to be 18

per cent for OECD countries, 38 per cent for transition countries, and 41 per cent for developing countries.

Price variability is common to developing country markets and chronic in agricultural commodity, input and labour markets. Variability is due to factors such as lack of credit, migration, inadequate infrastructure such as roads, irrigation and storage facilities, and erratic input supply. It is exacerbated by droughts, floods, pestilence, war and civil strife. If fairness motives influence decision making in response to price changes, then we would predict, on average, that resource allocation would be particularly affected in the more volatile markets of developing countries. Living rurally has also been a strong predictor of differences in time preference and risk tolerance, even when income is controlled for. Our prior is that these results may stem from the relatively high transaction costs faced by rural dwellers compared to their urban counterparts. In particular, the high cost of exchanging information, and thereby contextualizing and verifying it, may result in fewer, more homogenous, and less reliable or valid sources.

More individuals in developing countries live in rural areas, with South America being more urbanized than other regions. More than half of the population of most African countries live rurally. Two-thirds to three-quarters of the population live rurally in the largest countries in Asia (UNDP 2002). In the wealthier countries of North America and Europe, these numbers tend to be reversed. Rural populations also tend to be poorer and have access to fewer formal institutions and markets that might offer opportunities to insure against present consumption biases and perceived risks. Kinship and other informal means of solidarity tend to break down in periods of extreme hardship.

Transition economies vary widely in terms of rural living and poverty, but they share some characteristics such as limited market experience, a history of controlled information, and growing levels of inequality. In Russia, for instance, there is evidence that individuals measure their wellbeing relative to reference points, either their own baseline, or relative to others (Easterlin 1996; McBride 2001; Graham et al. 2004). For seven consecutive years after the 'fall of the wall' GDP per capita annual growth declined, and current levels remain below 1989 levels. The Gini coefficient for Russia in 2000 was 45.62, high by OECD country standards, higher than estimates of 29 in 1992, and most certainly higher than estimates prior to 1989 (Graham 2004: 6). Hence we can expect that some individuals, at least relative to their income pre-transition, perceive themselves as being worse off than previously.

Individuals in transition economies are facing vast changes in how scarce resources are allocated; prices and a market system are replacing time and personal networks. There is some evidence that suggests market experience is important to the distribution of rents and overcoming the endowment effect, but that with 'sufficient' experience learning occurs and competitive market outcomes can be expected (Camerer 1987; List 2003,

2004a, 2004b). The pace and extent of this learning, however, can be expected to vary with age and experience with the former regime.

It becomes rather obvious from the observations above that, in the context of developing and transition economies, policy based on standard economic model predictions may fail to achieve its objectives. This seems to be consistent with the evidence regarding the rather limited effectiveness of development assistance in producing the expected results both in terms of growth and poverty reduction. This poor performance has led to changes in the prevailing paradigms for development policy and the search for the proper balance between 'state-led' and 'market-led' solutions. We are suggesting that this distinction may in fact be artificial or, in any case, not the most important one. Rather, development policy frameworks may have missed salient points of peoples' behaviour which determine the success (or lack thereof) of the policies themselves.

3.1 How might behavioural economics contribute?

There are hundreds of models of decision making that more accurately describe behaviour in a particular context relative to strict rational maximization. But efficient policy requires less ad hoc behavioural models. Preserving the ability to generalize may come from retaining the rigor of the current model where its predictions have been regularly confirmed, and making changes only where behaviours systematically deviate for certain populations or in certain domains.

For example, the findings of time varying discount rates appear to be fairly robust across at least four cultures (Anderson et al. 2004). And experiments suggest that many of the poor are sophisticates who seek self-commitment mechanisms (Ashraf et al. 2004). Gugerty (2004) provides evidence from Kenya that women join rotating savings and credit associations because they help to provide a self-commitment mechanism. Yet much microfinance literature and practice still rests on the traditional assumption that individuals can only be made better off with access to credit since they can always chose not to avail themselves of it. A behavioural perspective would recognize that even with good intent, the initial borrowing and payback plan of individuals can change. As it stands, offering credit without acknowledging imperfect self-control, and hence without the appropriate institutional incentives, such as voluntary and mandatory savings, assumes away a possible source of repayment difficulties.

Empirically documenting regular deviations from standard rationality is the first step to identifying where the pay-offs from modifying the model, and hence the policy design, may be worthwhile. Experiments that reveal individuals are willing to suffer economic losses for reasons of fairness and revenge may help to explain why simply improving market or programme access for individuals does not guarantee participation. Interventions that affect the qualitative dimensions of risky outcomes, such as technologies that increase a farmer's sense of control even if they do not affect the magnitude of the risk, may alter the decision to adopt. Considering the endowment

effect in the design of land reform programmes may narrow the gap between buying and selling prices. When property rights are incomplete, the endowment effect might also explain an individual's willingness to expend more than the value of the right through violence, law suits, etc., in order to secure what they believe is already theirs (Mullainathan forthcoming: 18).

Ravi Kanbur's (2003: 2) observation that 'the development journals are by and large stuck in the rational choice paradigm', is as Knetsch notes ironically, predicted by the findings of behavioural economics (1995: 75). But the problem may run deeper than the status quo bias to which Knetsch refers and may include uncomfortable compromises to consumer sovereignty, 'ought we allow people to behave imprudently?' (Strotz 1956: 179). Camerer et al. (2003) have suggested 'asymmetric paternalism', regulations that would create large benefits for those who are not fully rational but impose little or no harm on those who are. But it is not always clear how to move from experimental findings to policy recommendations. Sample sizes are often small due to the complexity of the survey instruments, and not random, which limits our ability to generalize. And without more theoretical guidance, the empirical strategy is often more inductive than deductive, and more opportunistic than strategic.

In his 1993 Nobel Prize speech, Douglass North argued that explaining the performance of economies through time required a better understanding of institutions and cognition (North 1994). International development policy now reflects a growing understanding of the role of institutions, but despite Nobel prizes to Herbert Simon, Vernon Smith and the psychologist Daniel Kahneman, what we know about cognition has yet to regularly penetrate our analysis. Yet our research suggests that for the future of development policy, despite the challenges, getting the institutions right requires understanding patterns of imperfect rationality.

References

Ainslie, G. (1991). 'Derivation of "Rational" Economic Behaviour from Hyperbolic Discount Curves', *AEA Papers and Proceedings* 81: 334-40.

Allais, M. (1953). 'Le comportement de l'homme rationnel devant le risque, critique des postulats et axiomes de lécole Americaine', *Econometrica* 21: 503-46.

Anderson, C.L., A. Cullen, and K. Stamoulis (2006). 'Decision Making Along the Policy Chain in Vietnam', mimeo.

Anderson, C.L., and A. Cullen (2005). 'The Qualitative Dimensions of Risky Outcomes and Decision Making in Russia', mimeo.

Anderson C.L., M. Dietz, A. Gordon, and M. Klawitter (2004). 'Discount Rates in Vietnam', *Economic Development and Cultural Change* 52(4): 110-34.

Anderson, C.L., and N. Nevitte (forthcoming). 'Teach Your Children Well: Changing Values of Thrift and Saving', *Journal of Economic Psychology.*

Ashraf, N., D.S. Karlan, and W. Yin (2004). 'Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines', forthcoming in *Quarterly Journal of Economics* May 2006.

Benzion, U., A. Rapoport, and J. Yagil (1989). 'Discount Rates Inferred from Decisions: An Experimental Study', *Management Science* 35: 270-84.

Bereby-Meyer, Y., and M. Niederle. (2005). 'Fairness and Bargaining', *Journal of Economic Behaviour and Organization* 56: 173-186.

Binswanger, H. (1980). 'Attitudes Toward Risk: Experimental Measurement in Rural India', *American Agricultural Economics Association* 62: 395-406.

Camerer, C.F. (1987). 'Do Biases in Probability Judgement Matter in Markets? Experimental Evidence', *American Economic Review* 77(5): 981-97.

Camerer, C., and R. Hogarth (1999). 'The Effects of Financial Incentives in Economics Experiments: A Review and Capital-Labour-Production Framework', *Journal of Risk and Uncertainty* 19: 7-42.

Camerer, C.F., and R.H. Thaler (1995). 'Ultimatums, Dictators, and Manners', *Journal of Economic Perspectives* 9(2): 209-19.

Camerer, C., S. Issacharoff, G. Loewenstein, T. O'Donoghue, and M. Rabin (2003). 'Regulation for Conservatives: Behavioural Economics and the Case for 'Asymmetric Paternalism', *University of Pennsylvania Law Review* 151: 1211-54.

Cameron, T.A., and G.R. Gerdes (2002). 'Eliciting-Specific Discount Rates', unpublished manuscript.

Carver, J.W. (2003). 'Cross-Cultural Data Production: The Case of Russia', mimeo, Evans School of Public Affairs, University of Washington.

Chapman, G.B. (2003). 'Discounting of Health Outcomes, in Time and Decision', in G. Loewenstein, D. Read, and R.F. Baumeister (eds) *Time and Decision*, Russell Sage Foundation: New York.

Corso, P.S., J.K. Hammitt, and J.D. Graham (2001). 'Valuing Mortality-Risk Reduction: Using Visual Aids to Improve the Validity of Contingent Valuation', *Journal of Risk and Uncertainty* 23(2): 165-84.

Davies, E., and S.E.G. Lea (1995). 'Student Attitudes to Student Debt', *Journal of Economic Psychology* 16: 663-79.

Easterlin, R.A. (1996). *Growth Triumphant: The Twenty-first Century in Historical Perspective*, University of Michigan Press: Ann Arbor.

Fisher, I. (1930). *The Theory of Interest*, Macmillan: London.

Frederick, S., G. Loewenstein, and T. O'Donoghue. (2002). 'Time Discounting and Time Preference: A Critical Review', *Journal of Economic Literature* XL: 351-401

Graham, C., A. Eggers, and S. Sukhtankar (2004). 'Does Happiness Pay? An Exploration Based on Panel Data from Russia', *Journal of Economic Behaviour and Organization* 55: 319-42.

Graham, C. (2004). 'Can Happiness Research Contribute to Development Economics?' mimeo, The Brookings Institution: Washington DC.

Gugerty, M.K. (2004). 'You Can't Save Alone: Commitment in Rotating Savings and Credit Associations in Kenya', mimeo, University of Washington: Seattle.

Gurr, T.R. (1970). *Why Men Rebel*, Princeton University Press: New Jersey.

Harris, C., and D. Laibson (2001). 'Hypberbolic Discounting and Consumption', forthcoming in 8th World Congress of the Econometric Society.

Hausman, J. (1979). 'Individual Discount Rates and the Purchase and Utilization of Energy Using Durables', *Bell Journal of Economics* 10(1): 33-54.

Henrich, J., R. Boyd, S. Bowles, C. Camerer, E. Fehr, H. Gintis, R. McElreath, M. Alvard, A. Barr, J. Ensminger, N.S. Henrich, K. Hill, F. Gil-White, M. Gurven, F.W. Marlowe, J.Q. Patton, and D. Tracer (forthcoming). 'Economic Man in Cross-Cultural Perspective: Behavioural Experiments in 15 Small-Scale Societies', *Behavioural and Brain Sciences*.

Johnson, E.J., J. Hershey, J. Meszaros, and H. Kunreuther (1993). 'Framing, Probability Distortions, and Insurance Decisions', *Journal of Risk and Uncertainty* 7(1): 35-51.

Kahneman, D. (2003). 'Maps of Bounded Rationality: Psychology for Behavioural Economics', *American Economic Review* 93(5): 1449-75.

Kahneman, D., and A. Tversky (1979). 'Prospect Theory: An Analysis of Decision Under Risk', *Econometrica* 47: 263-91.

Kahneman, D., J. Knetsch, and R.H. Thaler (1986). 'Fairness as a Constraint on Profit Seeking: Entitlements in the Market', *American Economic Review* 76(4): 728-41.

Kahneman, D., J. Knetsch, and R.H. Thaler (1990). 'Experimental Effects of the Endowment Effect and the Coase Theorem', *Journal of Political Economy* 98(6): 1325-48.

Kahneman, D., P. Slovic, and A. Tversky (eds) (1982). *Judgement under Uncertainty: Heuristics and Biases*, Cambridge University Press: Cambridge.

Kanbur, R. (2003). 'Behavioural Development Economics', notes from the conference on Economics for an Imperfect World.

Knetsch, J.L. (1995). 'Assumptions, Behavioural Findings, and Policy Analysis', *Journal of Policy Analysis and Management* 14(1): 68-78.

Kreps, D.M. (1990). *A Course in Microeconomic Theory*, Princeton University Press: New Jersey.

Kuechler, M. (1998) 'The Survey Method', *American Behavioural Scientist* 42(2): 178-200.

Laibson, D. (1997). 'Golden Eggs and Hyperbolic Discounting', *Quarterly Journal of Economics* 112 (2): 443-78.

Lawrance, E. (1991). 'Poverty and the Rate of Time Preference: Evidence from Panel Data', *Journal of Political Economy* 99: 54-77.

List, J. (2003). 'Does Market Experience Eliminate Market Anomalies?', *Quarterly Journal of Economics* 118(1): 41-71.

List, J. (2004a). 'Neoclassical Theory Versus Prospect Theory: Evidence from the Marketplace', *Econometrica* 72(2) 615-25.

List, J. (2004b). 'Testing Neoclassical Competitive Theory in Multi-Lateral Decentralized Markets', *Journal of Political Economy* 112: 1131-56.

Loewenstein, G., and D. Prelec (1992). 'Anomalies in Intertemporal Choice: Evidence and Interpretation', *Quarterly Journal of Economics* 107: 573-97.

Loewenstein, G., and J. Elster (1992). *Choice Over Time*. Russell Sage Foundation: New York.

Machina, M.J. (1987). 'Choice Under Uncertainty: Problems Solved and Unsolved', *Journal of Economic Perspectives* (1): 121-54.

McBride, M. (2001). 'Relative-Income Effects on Subjective Wellbeing in the Cross-Section', *Journal of Economic Behaviour and Organization* 45: 251-78.

McNeil, B.J., S.G. Pauker, H.C. Sox, and A. Tversky. (1982). 'On the Elicitation of Preferences for Alternative Therapies', *New England Journal of Medicine* 306(21): 1259-62.

Mullainathan, S. (forthcoming). Development Economics through the Lens of Psychology, *Annual Bank Conference in Development Economics*.

North, D.C. (1994). 'Economic Performance Through Time', *The American Economic Review*. 84(3): 359-68.

Pate, E. (1983). 'Acceptable Decision Processes and Acceptable Risks in Public Sector Regulations', *IEEE Transactions on Systems, Man, and Cybernetics* SMC-13:113-124.

Prelec, D. (2001). 'Compound Invariant Weighting Functions in Prospect Theory', in D. Kahneman, and A. Tversky (eds) *Choices, Values, and Frames*, Russell Sage Foundation: New York.

Rabin, M. (1996). 'Psychology and Economics', mimeo.

Rabin, M. (1998). 'Psychology and Economics', *Journal of Economic Literature* 36(1): 11-46.

Rabin, M. (2002). 'A Perspective on Psychology and Economics', *European Economic Review* 46: 657-85.

Rubinstein, A. (1998). *Modeling Bounded Rationality*, MIT Press: Cambridge MA.

Samuelson, P. (1937). 'A Note on Measurement of Utility', *Review of Economic Studies* 4: 155–61.

Schneider, F., and R. Klinglmair (2004). 'Shadow Economies around the World: What Do We Know?', *Institute for the Study of Labour (IZA) Discussion Papers* 1043.

Slovic, P., B. Fischoff, and S. Lichtenstein (1979). 'Rating the Risks', *Environment* 21:14-20, 36-39.

Slovic, P., M. Finucane, E. Peters, and D.G. MacGregor (2002). 'The Affect Heuristic', in T. Gilovich, D. Griffin, and D. Kahneman (eds) *Heuristics and Biases: The Psychology Of Intuitive Thought'*, Cambridge University Press: New York.

Strotz, R. (1956). 'Myopia and Inconsistency in Dynamic Utility Maximization', *Review of Economic Studies* 23: 177.

Thaler, R. (1981). 'Some Empirical Evidence on Dynamic Inconsistency', *Economic Letters* 8: 201-07.

Thaler, R.H. (1991). *Quasi Rational Economics*, Russell Sage Foundation: New York

Tversky, A., and D. Kahneman (1974). 'Judgment Under Uncertainty: Heuristics and Biases', *Science* 185: 1124-31.

Tversky, A., and D. Kahneman (1983). 'Extensional Versus Intuitive Reasoning: The Conjunction Fallacy in Probability Judgement', *Psychological Review* 90: 293-315.

UNDP (2002). *Human Development Report*, Oxford University Press for the United Nations Development Programme: New York.

World Bank (2005). *World Development Indicators,* Oxford University Press for the World Bank: Washington DC.