

Psychology and Development Economics

Sendhil Mullainathan¹

(MIT and NBER)

June 2004

¹ I am grateful to Anne Case, George Loewenstein, Jim Mirrlees, Thomas Wang and an anonymous referee for comments. Parts of this paper are also clearly inspired by discussions about current work on related topics with Marianne Bertrand and Eldar Shafir, Abhijit Banerjee and Richard Thaler. This paper draws on a paper presented at the Annual Bank Conference on Development Economics at the World Bank.

Introduction

Economists often study scarcity, yet their conception of decision-making assumes an abundance of psychological resources. In the standard economic model, people are unbounded in their ability to think through problems. Regardless of complexity, they can costlessly figure out the optimal choice. They are also unbounded in their self-control able to costlessly implement and follow through on whatever plans they set out for themselves. Whether they want to save a certain amount of money each year or finish a paper on time, they face no internal barriers in accomplishing these goals. Furthermore, they are unbounded in their attention. They think through every single problem that comes at them and make a deliberative decision about each one. In this and many other ways, the economic model of human behavior ignores the actual bounds on choices (Mullainathan and Thaler 2001). Every decision is thoroughly contemplated, perfectly calculated and easily executed.

In contrast to the traditional model, a growing body of research interprets economic phenomena with a more modest view of human behavior. In this alternative conception, individuals are bounded in all of the above-mentioned dimensions (and more). Practically, this conception begins with the rich understanding of human behavior that experimental psychologists have developed through numerous lab and some field experiments. This view, ironically enough, emphasizes the richness of behavior that arises from scarcities and focuses on the bounds on cognitive ability, self-control, attention and self-interest. Theoretical models are now being constructed that help to incorporate these ideas into economic applications. Perhaps even more compelling is the recent empirical work that suggests the importance of these psychological insights for real behavior in contexts that economists care about. In a variety of areas, from asset pricing to savings behavior to legal decision-making, well-crafted empirical studies are challenging the traditional view of decision-making.

What does this research have to say for economic development? I begin by highlighting some areas where the existing research can be directly applied to development. The bulk of work on savings can be translated into understanding savings institutions and behavior in developing countries. Additionally, the insights about self-control have some direct links to understanding education, and the behavioral approach also appears to add some insight to the large body of research on the diffusion of innovation. The question of how (and when) to evaluate the impact of development policies can also be better understood. Yet since psychology is only beginning to make inroads into applied areas of economics, beyond these areas, few papers explicitly deal with psychology and development. I therefore speculate about additional specific areas where psychology could be useful in the future: poverty traps, conflict, social preferences and corruption, and research on the psychology of the poor.

An important caveat is in order. The attempt to incorporate psychology into development has a dubious history. In some cases, researchers have attempted to label the poor as especially “irrational” (or “myopic”) to explain their state. The research I describe here starts from a completely different presumption, that the poor have the same biases as everyone else.² It is neither an attempt to blame the poor for their poverty nor to argue that the poor have specific irrationalities. Instead my goal is to understand how problems in development might be driven by general psychological principles that operate for both poor and rich alike. When I speak of self-control, for example, I am speaking of those self-control problems that exist in equal measure around the world (e.g. the difficulty of resisting a tempting snack). These problems may matter more for the poor because of the context in which they live but their core is universal. (Bertrand, Shafir and Mullainathan 2003).

Immediate Barriers to Education

The rational choice model of schooling is straightforward (Becker 1993). Individuals trade off the costs and benefits of schooling to decide how much schooling to get. Benefits come in a variety of forms such as better jobs or better marriage prospects.

² This perspective has been put forward in Bertrand, Shafir and Mullainathan (2003).

Costs could be direct financial costs (e.g. fees) as well as any opportunity costs (e.g. foregone labor). In the case of children, of course, parents make the actual choices. They do so to maximize some combination of their own and their children's long run welfare, with their exact weight depending on their altruism.

This view of education ignores the richness of the hardships faced by parents trying to educate their children in a developing country. Consider an Indian village where a poor father Suresh is eager to send his son Laloo to school for the next school year. Suresh recognizes the value of his son's education in allowing Laloo to get a government job, marry better or simply live more comfortably in a rapidly changing world. To ensure that he has money for school fees, textbooks or perhaps a school uniform, Suresh begins to save early. Quickly he encounters some competing demands on the money. Suresh's mother falls ill and needs money to buy some anesthetics to ease her pain. Though his mother insists that her grandson's education is more important, he is torn. Enormous willpower is required to let his mother suffer while Suresh saves money that he knows could ease her pain. Knowing that he is doing what is best in the long run is small consolation in the moment. Suresh overcomes this hurdle and enrolls his son. After some weeks, Laloo starts to show disinterest. As with most children in the world, sitting in a classroom (and not an appealing one at that) is not very engaging, especially since some of his friends are outside playing. In the evenings, exhausted from tiring physical work and overwhelmed by the pressures of everyday life, how will he handle this extra stress? Will Suresh have the mental energy to convince his son of the value of education? Will he have the energy to follow up with the teacher or other students to see whether Laloo has actually been attending school? This fictional example illustrates merely one important tension in third world educational attainment. Even the best of intentions may be very hard to implement in practice, especially in the high stress settings that the poor inhabit.

These problems are intimately related to how people view future decisions in the present, a topic that psychologists and behavioral economists have studied extensively through experiments. Several experiments have resulted in theoretical principles used to

understand individual's decision making process, and these principles can provide a framework with which we can understand the difficulties the fictional Suresh faced in his attempts to implement his decision on his son's schooling.

Behavioral economists have recently begun to better understand the devices people may use to make decision across time with hyperbolic discounting theory. Generally, people value preferences differently across time and tend to exhibit short run impatience and long run patience. Formally, the theory is modeled by discount rates that vary with horizon. People have very a high discount rate for short horizons (decisions about now versus the future) but a very low one for distant horizons. Hyperbolic discounting is present in many facets of life as evidenced by the following examples.

Consider the following offer: would you rather receive \$15 today or \$16 in 1 month? More generally, how much money would I need to give you in one month to make you indifferent to receiving \$15 today? What about in 1 year? What about in 10 years? Thaler (1981) presented these questions to subjects and found median answers of \$20, \$50 and \$100. While at first blush these answers may seem reasonable, they actually imply huge discount rates that also vary hugely across time: 345% over one month, 120% over a 1-year horizon and 19% over a 10-year horizon.³ Subjects greatly prefer the present to the future.

These choices also imply that the extent of patience *changes* with the horizon. This is made most clear in the following choice problem:

Would you prefer \$100 today or \$110 tomorrow?

Would you prefer \$100 thirty days from now or \$110 thirty-one days from now?

³ One reason subjects may show such preferences is that they may doubt that they will actually get the money in the future, leading them to value it at a lower rate. While this may be an effect, the literature on discounting finds similar results even when these issues of trust are dealt with (Frederick, Loewenstein and O'Donoghue (2002)).

Many subjects give different answers to these two questions. In the first scenario, they often prefer the immediate reward (\$100 today), but in the latter, they often prefer the delayed reward (\$110 in thirty-one days).

Such preferences are inconsistent with the standard model. To see this, suppose people have a utility function u and discount the future at a daily rate δ . Then in scenario one, the value of \$100 today is $u(100)$ and the value of \$110 tomorrow is $\delta u(110)$. On the other hand, in problem two the value of the two options is $\delta^{30}u(100)$ and $\delta^{31}u(110)$. In simplified form, this is the exact same trade-off. In other words, with the standard constant discount rate, individuals should choose the same thing in both situations. The inconsistency with the two choices is resolved only if the discount rate is allowed to vary with time.

Differences in preferences for the immediate versus the future can also be seen in the field. Read, Loewenstein and Kalyanaraman (1999) ask subjects to pick three rental movies. The subjects either pick one by one for immediate consumption, or they pick all at once for the future. When picking sequentially for immediate consumption, they tend to pick “low-brow” movies. When picking simultaneously for future consumption, they tend to pick “high-brow” movies. Once again, when planning for the future, individuals are more willing to make choices that have long-run benefits (presumably “high-brow” movies) than when choosing in the present.

The difference in choices at different horizons poses a problem for the individual. Consider a concrete example. Suppose I face the chore of house cleaning. This weekend I am very busy and decide to delay it, but my preference for next weekend is that I should start my cleaning. What happens next week though? What was a decision about the distant future (where I exhibited patience) becomes a decision about the present (where I exhibit impatience). My choice may now change. Once again, the option of putting it off seems appealing, as appealing as it did last week when I made the same decision. In other words, there is a conflict between what I plan to do in the future and what I would actually do when the future arrives.

This conflict is one of the many difficulties parents face in getting their children educated. Returning to our fictional story, Suresh wanted his son Laloo educated and decided today that his future self would put in the effort and money to send Laloo to school. Yet in the moment, many immediate pressures impinge on his time, money and energy, making it hard for Suresh to implement his original plan. This view presumes that parents would like to see their children educated but simply can't find a credible way to stick to that plan.

This behavioral approach improves our understanding of many components of education. It provides explanation of the gap between parents' stated goals and actual outcomes. The Probe report on basic education in India finds that many parents are actually quite interested in education. Even in the poorest Indian states with the worst education systems, they find over 85% of the parents agreeing that it is important for children to be educated. In the same survey, 57% of parents feel that their sons should study "as far as possible". Another 39% feel that they should get at least a grade 10 or grade 12 education. Clearly parents in these regions value education. Yet this contrasts with very low educational attainment in these states. This gap is reminiscent of the gap between desired and actual retirement savings in the United States. In one survey, 76% of Americans believed that they should be saving more for retirement (Farkas and Johnson 1997). Though they want to save, many never succeed. As noted earlier, immediate pressures are even more powerful in the education context. Putting aside money to pay for schooling requires making costly immediate sacrifices. Fighting with children who are reluctant to go to school can be especially draining when there are so many other pressures. Walking a young child every day to a distant school requires a constant input of effort in the face of so many pressing tasks. Put another way, if a middle-class American, supported by so many institutions, cannot save as much as they want, how can we expect a Rajasthani parent to consistently and stoically make all the costly immediate sacrifices to implement their goal of educating their children?

This behavioral framework also helps to explain, in part, an interesting phenomenon in many developing countries: sporadic attendance. In contrast to a simple human capital model, education does not appear to follow a fixed stopping rule in which students go to school consistently until a particular grade. Instead, students go to school for some stretch of time, drop out for some stretch and then begin again. This sporadic attendance, though potentially far from optimal, is a consequence of the dynamically inconsistent valuation of preferences described earlier. When faced with particularly hard to resist immediate pressures, individuals will succumb to them. When these pressures ease, it once again becomes easier to implement the original plan of sending their child to school. More specifically, parents who have “slipped off the wagon” may find certain salient moments to refocus their plans for their child’s education. One empirical prediction here is that at the beginning of the school year, attendance should perhaps be higher than at any other time as many parents are newly inspired and decide to give it another try. As they succumb to immediate pressures, attendance would then decline throughout the year.⁴

One practical application of the behavioral perspective is that programs that make schooling more attractive to students may provide a low-cost way to make it easier for parents to send children to school. For example, a school meals program may make school attendance attractive to children and ease the pressure on parents to monitor their children’s attendance (see Vermeesch 2003 for a discussion of such programs). One could even be creative in coming up with these incentive programs. Offering school sports, candy, or other cheap inputs to children may have large effects on attendance. In fact, under this model, such programs could have extremely large benefit to cost ratios, much larger than could be justified by the monetary subsidy alone.

This perspective on schooling matches the complexity of life in developing countries. Of course, immediate pressures are not the only problem. Numerous other factors—from

⁴ This last point provides one potential way to distinguish this explanation from a rational model with very large liquidity shocks. Moreover, such a rational model faces difficulties if (1) parents rationally forecast liquidity shocks and (2) there are scale economies to attending for long continuous periods. In this scenario, parents would build a “buffer stock” early on to insure against such shocks and then send the child to school for one long (and presumably more productive) stretch.

liquidity constraints to teacher attendance—also play a role, but these areas are already well studied. The difficulty of self control, however, is not and deserves more scrutiny.

Demand for Commitment, Default Settings and Savings

The savings dilemma that, in part, drives the shortcomings in educational attainment is a more general issue in developing countries and links directly to behavioral issues that drive similar savings problems in developed nations. The difficulty of sticking with a course of action in the presence of immediate pressures has implications for how individuals save. In the standard economic model of savings, there is no room for such pressures. Instead, people calculate how much money is worth to them in the future, taking into account any difficulties they may have in borrowing and any shocks they may suffer. Based on these calculations, they make a contingent plan of how much to spend in each possible state. They then, as discussed before, implement this plan with no difficulty. Yet, for the poor in many developing countries, implementing such plans is much easier said than done. They face a variety of temptations that might derail their consumption goals.

Hyperbolic discounting is again a useful model with which we can understand individuals' actual saving behavior. As alluded to earlier, a key question in this model is whether people are sophisticated or naive in how they deal with their temporal inconsistency. Sophisticated people would recognize the inconsistency and (recursively) form dynamically consistent plans. In other words, they would only make plans on which they would follow through. Naïve people, on the other hand, would not recognize the intertemporal problem and would make plans today assuming that they'll stick with them, only to abandon them when the time comes. There are reasons to believe both views. On the one hand, individuals appear to consciously demand commitment devices, to help them commit to a particular path. On the other, they appear to have unrealistic plans. Perhaps the best fit of the evidence is that individuals partly recognize their time inconsistency.

The practically important component of this view is that the commitment implicit in institutions is crucial to understanding behavior.⁵ Institutions can help solve self-control problems by committing people to a particular path of behavior. Just as Ulysses ties himself to his ship's mast to listen to the song of the Sirens but not be lured out to sea by them, people rely on similar – albeit not so dramatic – commitment devices in everyday life. Many refer to their gym membership as a commitment device (“Paying that much money upfront every month really gets me to go to the gym”). Another example is Christmas clubs. While less common than before, they used to be a very powerful commitment tool for saving to buy Christmas gifts.

Suggestive evidence on the power of commitment devices is given in Gruber and Mullainathan (2002) which studies smoking behavior. Rational choice models of smoking treat it roughly like any other good. Smokers make rational choices about their smoking understanding the physiology of addiction that nicotine entails. Behavioral models on the other hand recognize a self-control problem in the decision to start smoking and in the decision (or rather attempts) to quit. There is some survey evidence suggestive of the behavioral model. Smokers often report that they would like to quit smoking but are unable to do so. This resembles the temporal pattern above. Looking into the future, smokers would choose to not smoke. But when the future arrives, they end up being unable to resist the lure of a cigarette today (perhaps promising that tomorrow they will quit). To differentiate these theories we examined the impact of cigarette taxes. Under the rational model, smokers are made worse off with these taxes. This is a standard dead-weight loss argument. Smokers who would like to smoke cannot now because of the higher price. In behavioral models with hyperbolic discounters, however, taxes could make smokers better off. The very same force—high prices driving smokers to quit—that is bad in the rational model is good in this model. Smokers who wanted to quit, but were unable to due to the inability to commit, are now better off. In the parlance of time-inconsistency models, the taxes serve as a commitment device.

⁵ Commitment devices might also play a role because of *inter*-personal rather than *intra*-personal problems. People may choose commitment as a way to avoid pressure from other family members or outsiders. The distinction between these two is quite important and further work is needed to distinguish between them.

To determine the effect of the cigarette taxes on smokers' welfare, we use self-reported happiness data. While such data are far from perfect, they can be useful in contexts such as this one.⁶ Using a panel of individuals in the United States and taking advantage of state variation in cigarette tax increases, we find that happiness of predicted smokers increases when cigarette taxes increase. Relative to the equivalent people in states with no rise in cigarette taxes (and relative to those who tend not to smoke in their own state), these predicted smokers show actual rises in self-reported well-being. In other words, contrary to the rational model and supportive of the behavioral model, cigarette taxes actually make those prone to smoke *better off*. This kind of effect is exactly the one I alluded to in the introduction: institutions (cigarette taxes in this case) have the potential to help solve problems within people as well as between people.

There is also evidence on people actively choosing commitment devices. Wertencroch (1998) argues that people forego quantity discounts on goods they'd be tempted to consume (e.g. cookies) in order to avoid temptation. This is a quantification of the often-repeated advice to dieters: don't have big bags of cookies at home. If you must buy tempting foods, buy small portions of it. Trope and Fischbach (2000) show how people strategically use penalties to spur unwanted actions. They examine people scheduled to undertake small, unpleasant medical procedures for which there's a potential time inconsistency problem. As the date of the procedure nears, people have a strong desire to skip the procedure. Trope and Fishbach show that people are aware of this problem and value a commitment device. They voluntarily agree to pay penalties for not undertaking the procedure, thus binding themselves to follow through with it. In fact, they smartly pick these penalties, by picking higher ones for more aversive procedures. Ariely and Wertebroch (2002) provide even more direct evidence. They examine whether people use deadlines as a self-control device and whether such deadlines actually work. Students in a class at MIT chose their own due dates for three different papers required in the course. The deadlines were binding so in the absence of self-control problems, they should clearly choose the latest due date possible for all the papers. They were told there

⁶ In this context, the variable of interest is relatively clean and thereby the mis-measurement is simply absorbed in the residual. See Bertrand and Mullainathan (2001) for a discussion of these issues.

were neither benefits for early due dates nor penalties for later ones, so they could only gain from the optional choice of being able to turn it in as late as possible. In contrast, students chose evenly spaced due dates for the three papers, presumably to give them incentives to get the papers done in a timely manner. Moreover, deadlines appeared to work. In a related study, they show that people given evenly spaced deadlines do better than those given one big deadline at the end.

Savings in developing countries can also be better understood through the behavioral perspective. It provides an alternative view on institutions such as roscas (Gugerty 2003). In a rosca, a group of people meets together at regular intervals. At each meeting, members contribute a pre-specified amount. The sum of those funds (the “pot” so to speak) is then given to one of the individuals. Eventually, each person in the rosca will get their turn and thus get back their contributions. Roscas are immensely popular but what is their attraction? They often pay no interest. In fact, given the potential for default (those who received the pot early may stop paying in), they may effectively pay a negative interest rate. One reason for their popularity may be that they serve as a commitment device in several ways. By making savings a public act, they allow social pressure from other rosca members to commit them to their desired savings level (Ardener and Burman 1995). As some rosca participants say, “you can’t save alone”. Each rosca member has incentives to make sure other members continue to contribute. Given self-control limitations, they also allow individuals to save up to larger amounts than they normally could achieve. Imagine someone who wished to make a durables purchase (or pay school fees) of 1000 rupees. By saving alone and putting aside money each month, they face a growing temptation. When they reach 400 rupees, might not some other purchase or immediate demand appear more attractive? The rosca doesn’t allow this temptation to interfere. Individuals get either nothing or the full 1000 at once. This all or nothing property might make it easier to save enough to make large purchases. It also helps to provide a more nuanced view of individual’s demand for liquidity. In the standard logic, the poor unconditionally value liquidity. After all, liquidity allows people to free up cash to attend to immediate needs that arise. If a child gets sick, money is needed to pay for medicine. This might be especially true for the poor. Shocks that are

small for the well off can be considerable for the poor and force them to dip into real savings. The poor, in these models, however face a tradeoff. They value liquidity for the reasons cited above but liquidity for them is also a curse since it gives them too easy access to savings. Consequently, durable goods and illiquid savings vehicles may actually be preferred to liquid savings vehicles. Liquid holdings like cash are far too tempting and spent too readily. On the other hand, by holding their wealth in items such as jewelry, livestock, and grain, individuals may effectively commit themselves to savings and resist immediate consumption pressures. In these models, therefore, there is an optimal amount of liquidity. Even when liquidity is provided at zero cost, they will choose some mix of illiquid and liquid assets.

Another implication of the behavioral perspective is that revealed preference arguments do not necessarily hold; hence, measuring policy success can be difficult. Observing that people borrow at a given rate (and pay it back) does not necessarily mean that the loan helps them; we must also examine the use of the loan. In some cases, it clearly benefits borrowers by helping them deal with a liquidity shock. In other cases, the borrower may use the loan to satisfy immediate temptations and end up straddled with debts. This distinction is important for understanding micro-credit in developing countries. Often, the metric of success for such programs is whether programs are self-sustainable, but this measure is valid only in the standard model where revealed preference applies. In this model, profitability implies that people prefer getting these loans even at a non-subsidized rate; revealed preference then implies their social efficiency. Yet in the presence of time inconsistency, the profitability and social efficiency of micro-credit potentially have little relationship. The key question is to what extent the loans exaggerate short-run impatience and to what extent they solve long-run liquidity constraints.⁷ Ultimately one needs a deeper understanding of what drives borrowers. One source for this might be data on loan usage. Are loans being spent on long-run investments (as is often touted) or spent on short-run consumption? Of course, some short-run consumption might well be efficient, but this data, combined with an

⁷ To make this contrast stark, note that in the United States, payday loan companies are a very profitable form of micro-credit.

understanding of the institution, would help to better understand (and improve) the social efficiency of micro-credit.

Policy can also provide cheaper and more efficient commitment devices. Saving in grain is an effective, but extremely risky, commitment device. Vermin may eat the grain or the interest rate earned on the grain could be zero or even negative. Moreover, it is important to recognize, that even if people demand such commitment devices, the free market may not do enough to provide them. The highly regulated financial markets in developing countries may prevent adequate innovation on these dimensions. Monopoly power may also lead to inefficient provision of these commitment devices if a monopolistic financial institution can extract more profits by catering to consumption temptations over the desire for commitment devices. In this context, governments, NGOs and donor institutions can play a large role by promoting such commitment devices.

Ashraf, Karlan and Yin (2004) give a stunning illustration of the role of commitment devices in savings. They offer depositors at a bank in the Philippines the opportunity to participate in “SEED” accounts. These accounts are like deposit accounts, except individuals cannot withdraw deposits at will. Instead, they can only withdraw the money at a prespecified date or once a prespecified savings goal has been reached. This account does not pay extra interest and is illiquid. In most economic models, people should turn down this offer in favor of the regular accounts offered by the bank. Yet they find large demand for them. More than 30% of those offered the accounts take them up. They also find these accounts help individuals to save. Six months later, those offered the accounts show substantially greater savings rates. Experiments such as these will contribute to a deeper understanding of savings and greatly improve development policy.

Defaults and Financial Institutions

Financial institutions do not simply help savings through their commitment value. A very important set of results in behavioral economics suggests that institutions also affect behavior simply through the defaults they produce. Madrian and Shea (2001) conducted a particularly telling study along these lines. They studied a firm that altered the choice context for employee participation in their retirement plan. When employees join the

firm, they are given a form that they must fill out in order to participate in the savings plan. Though the plan is quite lucrative, participation was low. Standard economic models might suggest that the subsidy ought to be raised. This firm instead changed a very simple feature of its program. Prior to the change, new employees received a form that said something to the effect of “Check this box if you would like to participate in a 401(k). Indicate how much you’d like to contribute”. After the change, however, new employees received a form that said something to the effect of “Check this box if you would like **not** to have 3% of your pay check put into a 401(k)”. By standard reasoning, this change should have little effect on contribution rates. How hard is it to check off a box? In practice, the study shows a large effect. When the default option is to not contribute, only 38% of those offered contributed. When the default option was contribution, 86% contributed. Moreover, even several years later those exposed to the contribution default still show much higher contribution rates.

While we cannot be sure from this data what people are thinking, I would speculate that some combination of procrastination and passivity played a role. Surely many looked at this form and thought, “I’ll decide this later”. But later never came. Perhaps they were distracted by something more interesting than retirement planning. In any case, whatever the default on the form was, they ended up with it. In fact, as other psychology tells us, as time went on they may well have justified their “decision” to themselves by saying, “3% is what I wanted anyway” or in the case of the “not to contribute” default, “that 401(k) plan wasn’t so attractive. In this way, their passivity made their decision for them. By making the small active choice to choose later, they ended up making a large decision about thousands of dollars of money.

These insights can also help us design new institutions. One example is Save More Tomorrow, a program created by Thaler and Benartzi (2003). The basic idea of Save More Tomorrow is to get people to make one active choice but to have them make it in such a way that if they remain passive afterwards they are still saving. Participants (who have chosen to save) decide on a target savings level and agree to start deductions at a small level from their paycheck *next year*. Each year thereafter, as they receive a raise in their income, their deductions increase until they hit their personal target savings level.

They can opt out of the program at any time, but the cleverness of the program is that if they do nothing and remain passive, they will continue to save (and even increase their savings rate).

The results have been stunning. In one firm, for example, more than 75% of those offered the plan participated in it rather than simply trying to save on their own. Of these, less than 20% opted out of the program. As a result, savings rates rose sharply. By the third pay increase (as the default increases cumulated), individuals had more than tripled their savings rates. But perhaps the greatest success has been the diffusion of this product. Many major firms and pension fund providers are thinking of adopting it and participation in the program will likely soon number in the millions. Save More Tomorrow is an excellent example of what psychologically smart institutional design might look like in the future. It does not solve a problem between people but instead helps solve a problem within the person: not saving as much as they would like.⁸

One simple but powerful implication of these results is that behavior should not be confused with dispositions (Bertrand, Shafir and Mullainathan 2004). An economist observing the savings behavior of a middle-class American and a rural farmer might be tempted to conclude something about different discount rates. In the standard model, the high savings of the middle-class American surely indicates greater patience. This need not be the case. Such an inference could be just as wrong as inferring that those who defaulted into 401(k) are more patient than those who were not defaulted into it. In other words, the behavioral difference between the farmer and the middle-class American may be a consequence of better institutions facilitating more automatic, default savings.

Default behavior has potentially large implications for improving savings in the developing world through bank reform. The lessons learned in the United States could easily be transferred to parts of developing countries. Institutions such as direct deposit of paychecks (into a savings account) could be very powerful in spurring savings. This banking innovation could be a very inexpensive and far-reaching device to improve savings rates of the middle-class in developing countries.

⁸ In this short space, I cannot do justice to all the psychological tools that the SMarT plan relies on. The full discussion in the original paper is well worth reading as an example of how to use psychological tools to better design policy.

On a more basic level, the evidence on default behavior implies that the introduction of basic banking to rural areas could in and of itself have large impacts on behavior. While deposit accounts are not as powerful of a default as direct deposit, a basic deposit account, as opposed to cash in hand, has a psychological element of illiquidity. The person then has to make one active decision—putting the money into the account—but then the act of saving the money becomes a passive one. When cash is available in the household, an individual must take action to save it. When money is in the bank account, one must take action to spend it. In this sense, a bank account may serve as a very basic commitment device. By keeping the money at a (slight) distance, spending it may be a lot less tempting.

Status Quo Bias and Diffusion of Innovations

Another broad area of significance to development is the diffusion of innovation. A large literature has focused on the slow rate of adoption of new technologies (Rogers 2003). This literature has been especially important in developing countries where a long-standing puzzle is why innovations, ranging from the Green Revolution to new medical technologies to fertility practices, appear to diffuse slowly.⁹

Several stylized facts emerge from this largely empirical literature. First, innovations generally follow an s-shaped diffusion, starting off slowly, speeding up suddenly (after a “tipping point”) and then slowing down to a steady state level.¹⁰

Second, adopters can be classified according to the timing of their adoption. The first adopters called “innovators” are only weakly connected in a social network sense. The second group of adopters, or “early adopters”, is comprised of those who are best connected to the rest of the network. They are often opinion leaders, enjoy high status

⁹ See Besley and Case (1994), Foster and Rosenzweig (1996), Munshi (2003) and Conley and Udry (2004) for empirical work in development economics on diffusion of agricultural innovations. The work on fertility is well summarized in Retherford and Palmore (1983) and Montgomery and Casterline (1998).

¹⁰ While much has been made of the shape of this diffusion process, it is unclear why this fact is so surprising. Many processes that start at zero and end at some positive diffusion would have this “s-shape”.

and better-educated. After this sub-classification adopts, the rest of the group follows. The final adopters, or “laggards”, are socially isolated like the “innovators”.

Third, diffusion takes place along social network lines. The time series evidence on adoption clearly supports this view. Slightly more direct evidence can be found in qualitative surveys (and in some quantitative, subjective ones) that ask people what drove their adoption decision. In a few cases, very compelling, empirical evidence exists. An early family planning experiment in Taiwan in 1964 showed social networking effects (Palmore and Freedman 1969; Lu et al. 1967). In this randomized experiment, a direct family planning intervention targeted at a specific group of subjects was shown to have some impact in the targeted group’s behavior. But most interestingly, those who were not directly contacted by the program but socially connected to the treatment group also showed effects.

A theoretical implication of these findings is that there are multiplier effects in policy interventions. If behavior-changing policies target some people, especially opinion leaders, then the interventions may also have an effect on the broader population. Montgomery and Casterline (1993, 1998) illustrate the multiplier effect in their study of a Taiwanese family planning program. Ignoring diffusion (and hence the multiplier effect), the study finds that program efforts accounted for no more than 5 to 20 percent of marital fertility decline. However, taking into account social diffusion increased the estimates of the program effect to account for over 30 percent of the decline. Multiplier effects have also influenced the shape of policy. In many cases, programs have enlisted specific local groups and institutions. Agricultural extension and doorstep delivery of contraception have all been modified with an eye towards utilizing multiplier effects.

While this work has been quite interesting, the greater incorporation of psychology into this social diffusion literature might increase our understanding in several ways. Much of the slowness of diffusion may come from status quo bias. Samuelson and Zeckhauser (1988) document a variety of phenomena known as the status quo bias, a principle most clearly shown by example. A group of subjects is given the following choice:

You are a serious reader of the financial pages but until recently have had few funds to invest. That is when you inherit a large sum of money from your great uncle. You are considering different portfolios. Your choices are:

Invest in moderate-risk company A. Over a year's time, the stock has a .5 chance of increasing 30% in value, a .2 chance of being unchanged, and a .3 chance of declining 20% in value.

Invest in high-risk company B. Over a year's time, the stock has a .4 chance of doubling in value, a .3 chance of being unchanged, and a .3 chance of declining 40% in value.

Invest in treasury bills. Over a year's time, these will yield a nearly certain return of 9%

Invest in municipal bonds. Over a year's time, they will yield a tax-free return of 6%.

A second set of subjects is given the same choice but with one small difference. They are told that they are inheriting a portfolio from their uncle in which most of the portfolio is currently invested in moderate-risk company A. The choice now is subtly changed. It is how much of the portfolio to *change* to the other options above. Interestingly, they find a large difference between the two treatments: much more of the money is re-invested in A when it is the status quo. This bias towards the status quo appears to run quite deep and is not just due to some superficial explanations (such as information content of the uncle's investments). Samuelson and Zeckhauser (1988) demonstrate this with a very interesting piece of evidence from the field. In the 1980s, Harvard University added several plans to its choice of health plans. This provides an interesting test of status quo bias: how many of the old faculty chose the new plans and how many of the newly joined faculty chose the old plan? They find a stark difference. Existing employees "chose" the older plans at a rate two to four times higher. In other words, incumbent employees make

the easiest choice of all: to do nothing. This bias towards the status quo creates a natural slowness of adoption. Not only must the innovation overcome beliefs that it is not worthwhile, it must also overcome the natural inertia intrinsic in behavior.

There is also a large literature in social psychology, which shows the power of social pressure towards conformity (see Cialdini and Trost 1998 for a review). Muzafer Sherif demonstrated the role of social pressure in a classic psychophysics study (Sherif 1937). Subjects were seated in a totally dark room facing a pinpoint of light some distance from them. After some time where nothing happens, the light appears to “move” and then disappear. Shortly thereafter, a new point of light appears. It too moves after some time and then disappears. Interestingly, this movement of the light is a pure psychophysical phenomenon known as the autokinetic effect. The light does not actually move; the eye merely makes it appear to move. The subjects were put in this context for repeated trials (many different resets of the light) and asked to estimate how far it had “moved”. When performed alone, these estimates were variable, ranging from an inch to several feet. However, an interesting pattern developed when subjects did this task in groups of two or three. Subjects’ estimates invariably began to converge on a particular number. A group norm quickly developed. In one variant, a member of the group was a confederate, someone who worked for the experimenter and gave a specific number. He found that the group quickly converged to the confederate’s answers. Others have found that norms manipulated in this way persist for quite some time. Even when subjects are brought in up to a year later, they show adherence to that initial norm. In the context of the Sherif experiment design, Jacobs and Campbell (1961) have shown how norms can be transmitted across “generations” of subjects. Suppose subjects 1 and 2 initially converge to a norm, then subject 1 is replaced by 3 for enough trials and then 2, by 4. The final group consisting of totally new subjects 3 and 4 will conform to the norm previously established by 1 and 2.¹¹

¹¹ Camerer and Weber (2003) present an interesting examination of how such norms can arise and evolve over time.

Solomon Asch expanded these results to an even simpler task. Subjects enter a lab and sit with others (confederates) to judge the length of lines such as those in Figure 1. The subject hears the judgment of the others and then makes his own. For several trials, this is a very boring task as it is pretty obvious which line is longer. But then there is a twist. On one of the trials, the first person makes a wrong choice. Then the second person makes the same wrong choice. And so it continues until it is the subject's turn. What Asch found was stunning. Between 50 to 80 percent of the subjects yielded to the erroneous majority at least once. Of course, as Asch notes, it is not the subjects' perception of the line length (unlike what is perhaps the case in the Sherif experiment) that is altered. Many people are simply willing to conform in their behavior.

Similar effects have been found in more naturalistic settings. For example, numerous studies have used high-frequency data to argue for contagion effects in suicides. Media reporting of a suicide often leads to more suicides in areas covered by the newspaper in the days immediately following the reporting (Phillips 1979 and 1980). While information stories could potentially be told, a social imitation story is much more plausible for these findings

These and numerous other experiments highlight the sheer power of social pressure to drive behavior. The traditional economic model acknowledges imitative behavior across networks but explains the phenomenon with an information diffusion story. However, some of the social network effects that drive slow diffusion of innovation may also be traced to social pressures and conformity. An important question in this literature is to empirically disentangle these two effects – information diffusion and social conformity – and to theoretically understand the ways in which they differ.

Even when information diffusion is the explanation of network effects, psychology can be helpful in understanding the impact of information. In particular, most models of diffusion (including the Bayesian model) crucially assume that the most useful information diffuses and has the correct impact on beliefs. This need not be the case. For example, experimental studies have shown that the salience of a piece of data, independent of its information content, affects how it will influence beliefs. Recent experimental work is beginning to examine what data people choose to pass on to their

social network. The evidence suggests that emotional factors can have large effects, once again independent of the information content (Heath, Bell and Sternberg 2001). These results could be useful in interpreting diffusion, especially determining which ideas tend to diffuse well (those with salient positive outcomes) and which tend to diffuse poorly (those with non-salient benefits or with salient costs).

For example, a recent study by Miguel and Kremer (2003) on social learning about deworming pills can be fruitfully interpreted in this light. They selected several schools at random and offered them deworming drugs. Using experimental variation, they estimated that the deworming drugs had extremely large and positive health and education benefits. Miguel and Kremer further exploited this randomization to see the impact on non-selected schools in the following year when they too were offered the opportunity to take the deworming pills. They examined children who had friends in the treatment schools (and hence had friends who had already had experience with the pills). Given the positive benefits of the drugs, social learning models would imply these children would be much more likely to take the deworming pills. In contrast, Miguel and Kremer find that they are much less likely to take the drug. This “negative” social influence effect can be understood in terms of salience.¹² Deworming pills, while beneficial in the long-run, have short-run costs that are particularly salient. Children can have stomach cramps, vomiting and other side effects as the worms work their way out of the system. This asymmetry in salience of the costs and benefits suggests that the costs may well spread more than the benefits (and be more influential on choice). On average, children and parents may hear more stories about the short term side effects than the long-run benefits.

Slow diffusion may also be a consequence of the earlier discussed self-control bounds that hamper efforts to save. People may be interested in adoption but still have delays in getting the required savings or undertaking the other actions required to adopt. This is made most transparent in recent work by Duflo, Kremer and Robinson (2004). They examine the decision by Kenyan farmers to adopt fertilizer. The fertilizer clearly shows

¹² Miguel and Kremer themselves prefer a standard economic interpretation in which expectations of the deworming pills were even higher than their realized benefits.

large positive productivity benefits, but they find that even when farmers can see its benefits, they do not necessarily adopt it. Even in cases where they adopt it, they often revert to old practices after a few years. Using a set of innovative financial interventions, Duflo, Kremer and Robinson, argue that the psychology of financial management plays a role. It appears that though the farmers would like to adopt the fertilizer they are unable to save up enough money to do so. Interventions that help them to overcome this savings problem appear to greatly increase adoption.

As a whole the literature on diffusion of innovations appears to be going through a period of change. Insights from psychology are helping to uncover new phenomena as well as to cast new light on old ones.

Self Serving Bias and Evaluation

Evaluation is one of the important issues facing development practice. When should a particular development program be evaluated? How should it be evaluated? Psychology may help us better understand the social process by which these questions are answered. An often-repeated finding in psychology is that beliefs and the perceptions that feed into forming them can be biased. This bias may be the result of prior beliefs, personal affiliations, or a desire for a particular outcome. Hastorf and Cantril (1954) got two groups of students, one from Princeton and one from Dartmouth, to watch film of a Princeton-Dartmouth football game. Each student was asked to count the number of penalties committed by both teams as determined by their personal judgment (and not the official count from game referees). Though both groups watched the exact same tape, the counts show that they “saw a different game”. Dartmouth students saw an equal number of flagrant and mild penalties committed by both teams. Princeton students, on the other hand, counted three times as many flagrant penalties by Dartmouth than flagrant penalties by Princeton and the same number of mild penalties. Clearly, school affiliations influenced the game they “saw”.

Babcock and Loewenstein (1997) provide a particularly stunning example of the self-serving bias. Subjects are asked to bargain over how to deal with a particular tort case

that is based on a real trial that occurred in Texas. Researchers assign each subject the role of lawyer for the defendant or plaintiff. They read all the case materials and then bargain with each other over a settlement amount. If they fail to settle, the award will be what the judge decided in the actual case (which is unknown to the subjects at the time of bargaining). Interestingly, subjects are paid as a function of the final settlement amount and pay an explicit additional cost if they must go to the judge for a decision. Subjects are also asked to assess (in private) how much they think the judge will award them. Finally, some pairs of subjects read the entire description of the case *before* knowing what role they'll play. Others read it afterwards. This order of reading has a large effect. Those who read the case prior to their role assignment settle at a 94% rate and don't require a decision by the judge. But those who read afterwards only settle at a 72% rate. Moreover, as a rule, those who were assigned their roles prior to reading the case details tended to exaggerate how much the judge would favor them. In short, they exhibit quite biased beliefs based on their role. These conflicting beliefs are generated through nothing more than the assigned roles. Even more significant is that this role-playing bias goes against their real life material interests in one important way. They must pay to go to court yet their biased beliefs force them to do so much more often. Much like with the football game, these subjects saw very different cases. In some sense, they saw what they "wanted" to see.

These examples of self-serving bias have far-reaching implications for all development policy and underscore the importance of randomized approaches. Implementers of specific policies will largely have this natural, self-serving bias, which may cloud the evaluation of an intervention's success. A useful example of this danger is the study of Cabot's intervention program for delinquent youth in Cambridge and Somerville, Massachusetts (Powers and Whitmer 1951). This intervention combined all the best available tools at the time for helping these delinquent youth, from tutoring and psychiatric attention to interventions in family conflicts. Those involved in the program raved of its success. They all had very positive impressions. What made the program unique however was that a true random assignment procedure was used to assign the

students. When this data was examined, contrary to the very positive (and likely heartfelt impressions of the caseworkers), there was little program effect.

Ross and Nisbett (1991) cite another interesting example of self-serving bias: a meta-analysis by Grace, Muench and Chalmers (1966) who study all medical research on the “portacaval shunt”. This was a popular treatment for cirrhosis of the liver and 51 studies examined the efficacy of this treatment. The doctors and scientists doing these studies all had the same good intention: to see whether this procedure worked. But the studies differed in one important way: 15 of them used controls but without randomization while 4 of them used truly randomized strategies. Thirteen of the fifteen non-randomized studies were “markedly or moderately enthusiastic” about the procedure. Yet only one of the randomized studies was “markedly or moderately enthusiastic”.

What is going on here? I believe the good intentions of the doctors and scientists compromise the evaluation. There is always subjectivity in non-randomized trials – what controls to include, what controls not to include, which specification to run, etc. Such subjectivity gives room for self-serving bias to rear its head. Exactly because the researchers on these topics are well-intentioned and exactly because they hope the procedure works, it is all too easy for them to find a positive result. Much as with the Dartmouth and Princeton students, they saw in some sense what they wanted to see.

These examples highlight the damaging role of self-serving bias in evaluation. Especially in the context of development where most people working on a project would like to see it succeed, it is all too easy for self-serving bias to affect evaluations. Beyond the obvious econometric benefits of randomized evaluation, the reduction of researcher’s latent biases is one of randomization’s greatest practical benefits. The methodology allows us to escape the dangers of biased perception, of which researchers or field workers are no freer than anyone else in the population. One of the big changes in modern development research is the greater reliance on randomized trials. This will better insulate us from the psychological fallibilities of researchers and policy makers.

Interesting Further Directions

Self Control and Poverty Traps

The evidence on self-control bounds suggests that savings can be psychologically difficult. Can these results help in theoretical models of poverty traps? In ongoing work with Abhijit Banerjee, I am investigating this question. Traditional models of poverty traps often rely on investment non-convexities. If investment projects require, for example, a certain discrete amount to earn extremely large returns, then there can be a reinforcing nature to poverty. With liquidity constraints, the poor may not be able to save enough to avail themselves of this investment opportunity.

Self-control problems present the possibility of a different kind of poverty trap. Under the assumption that temptations are concave in income, richer individuals face fewer temptations, proportionally, than poorer ones.¹³ Suitably modeled, such temptations could in principle generate poverty traps themselves. The poor suffer from low savings rates since temptations effectively “tax” a greater proportion of income. The rich on the other hand, while they face and succumb to similar temptations, are not as affected in a proportional sense. This can make poverty an absorbing state. Understanding the theoretical underpinnings and empirical importance of this type of model could be a fruitful area for future work.

Loss Aversion and Property Rights

Consider the following simple experiment. Half a room of students is given mugs and the other half is given nothing (or a small cash payment roughly equivalent to the value of the mugs). The subjects are then placed in a simulated market where a mechanism determines an aggregate price at which the market clears. How many mugs should change hands? Efficiency dictates that market clearing should allocate the mugs to the 50% of the class that values it the most. Since the mugs were initially randomly assigned, trading should have resulted in exactly half the mugs changing hands. Kahneman,

¹³ For example, temptations such as especially tasty desserts are sub-linear. The presence of some temptations like those whose cost do not scale with income would create a declining importance of temptations as a proportion of income.

Knetsch and Thaler (1990) have in fact run this experiment. Contrary to the simple prediction, however, they find a stunningly low number of transactions. Roughly 15% of the mugs trade hands. This oddity is explained by a phenomenon called the *endowment effect*, in which those given objects appear to very quickly come to value them more than those not given the object. In fact, those who were given mugs had a reservation price three *times* that of those who were not given the mug.

This phenomenon reflects in part a deeper fact about utility functions: prospect theory, where people's utility functions are defined in large part on changes. In the traditional model of utility people would value the mug at $u(c+Mug)-u(c)$. Notice the symmetry: both those with and without the mug value it the same (on average). Under prospect theory, utility is defined by a value function that is evaluated locally and in changes. Those who receive the mug consider its loss as a function of $v(-Mug)-v(0)$. Those who do not receive the mug value its gain at $v(Mug)-v(0)$. In this formulation, valuations are asymmetric. The difference in valuation between the two depends on whether $v(Mug)$ is bigger or smaller than $-v(-Mug)$. The evidence above is consistent with a variety of evidence from other contexts: losses are felt more sharply than equivalent gains. Thus $v(x) < -v(-x)$. This phenomenon, known as loss aversion, has been seen in many contexts.¹⁴

The insight about loss aversion can also help us understand why policy change is so difficult in developing countries. Consider market reforms that transfer resources from one group to another with an efficiency gain. Suppose privatizing a firm will result in gains for customers while resulting in losses for incumbent workers. Prospect theory helps to explain why such reforms are fought so vigorously – the losses are felt far more sharply by the workers. One implication of loss aversion is, at the margin, to pursue strategies that preserve the rents of incumbents rather than ones that try to buy out incumbents. All other things equal, a strategy that offers a buyout for incumbent workers will be far more costly than one that grandfather them in. The buyout requires the government to compensate them for their loss and this can be much larger than simple

¹⁴ See Odean (1998) and Genesove and Mayer (2001) for two good examples.

utility calculations can suggest. In contrast a strategy that guarantees incumbent workers a measure of job security would not need to pay this cost.¹⁵ Many situations of institutional change require some form of redistribution. The recognition of loss aversion suggests successful policies may require more consideration of the losses of incumbents.

Loss aversion also reinforces the importance of well-defined property rights. Consider a situation where there is a single good, such as a piece of land L . There are two individuals, A and B, who can engage in force to acquire (or to protect) the land and engaging in violence may result in acquisition. In the presence of well-defined property rights (say this land belongs to person A), the decision to engage in force is straightforward. If B engages in force he stands to gain $v(L)$ if his force is successful. A on the other hand stands to lose $v(-L)$ if he doesn't engage in force. In this case, loss aversion implies that A stands to lose a lot more than B could gain. So with well-defined property rights A would engage in more force than B. Consequently, B may never attempt force. Even in the absence of enforcement, loss aversion may mean that well-defined property rights may deter violence.

Consider now the case of ill-defined property rights. Suppose that both are unsure who owns the land. Specifically take the case where they both think they own it. This is an approximation to the situation where ownership with probability $\frac{1}{2}$ already gives a partial endowment effect, or to the situation of biased beliefs where both parties may have probability greater than $\frac{1}{2}$ of owning it. In this case, both A and B think they stand to lose $v(-L)$ if they do not fight for the land. In other words, in the absence of well-defined property rights, both parties will put in large amounts of resources to secure what they already think is theirs. This is one of the powerful implications of loss aversion. Appropriately defining property rights prevents two (or more) parties from having an endowment effect on the same object. Conflicting endowments such as this are sure to produce costly attempts at protecting the perceived endowments. Anything from costly territorial activities (fencing and de-fencing) all the way to violence may result.

Deleted: to me

¹⁵ Of course, this is a comparative static only. In any given context, there may be pressing reasons to favor one policy over the other.

Social Preferences and Corruption

In many important development contexts, self-interested behavior is very deleterious. Bureaucrats in many countries are corrupt. They enforce regulations sporadically or take bribes. Another stark example is teacher absenteeism. Numerous studies have found that teacher absenteeism is one of the primary problems of education in developing countries. Teachers simply do not show up to school, and as a result, little education can take place. This blatantly selfish behavior stands in contrast to some evidence on social preferences that individuals may value the utility of others. I will review this literature and describe how social preferences may be contributing to the problem and may serve as part of the solution.

A very simple game called the ultimatum game has become an excellent tool for studying social preferences (Guth, Schmittberger and Schwarz 1982, Thaler 1988). In this game, one player (call him the “Proposer”) makes the first move and offers a split of a certain amount, say \$10. The second player (“Responder”) decides whether to accept or reject this split. If it is accepted, P and R get the proposed split. If rejected, then both get zero. This game has been run in many, many countries and for stakes that range from a few dollars in the US to a few months of income in many countries. Yet the pattern of findings is relatively constant.¹⁶ First, responders often reject unfair offers (i.e. away from 50-50 splits). Second, proposers often make very fair offers, close to 50-50 or 60-40. Moreover, proposers’ fair offers are not just driven by fear of rejection. They tend to make offers larger than what a simple (risk-neutral) fear of rejection implies.¹⁷ The ultimatum game illustrates two facts about interpersonal preferences: (i) people care about others and are willing to give up resources to help others and (ii) people react negatively to *perceived* unfair behavior and are willing to give up resources to punish it. The second fact illustrates part of the “dark side” of interpersonal preferences. In simple altruistic models, interpersonal preferences are only a good thing: having one person care

¹⁶See Heinrich et. al., (2002), however, for some interesting differences in some tribal cultures.

¹⁷ This is most directly seen in a variant of the Ultimatum game, called the dictator game. Here the Proposer makes an “offer” but the Responder has no choice but to accept it. In this game, the threat of rejection is removed and one continues to find non-zero offers by the Proposer, though they are lower than in the Ultimatum game.

in a positive way about another only makes it easier to deal with externalities and so on. The Responder's possible "punishment" behavior shows, however, the way in which interpersonal preferences could potentially cause inefficiencies and conflicts.

This possibility is clearest in a classic experiment by Messick and Sentis (1979). They ask subjects to imagine they have completed a job with a partner. They are asked to decide what "fair" pay for their work is. They divide the subjects into two groups however. One group is told to imagine that they had worked 7 hours on the task while the partner had worked 10. The other group is told to imagine that they had worked 10 hours while the partner had worked 7. Both groups were told that the person who had worked 7 hours had been paid \$25 and were asked what the 10 hour person should be paid. Those who were told that they had worked 7 hours (and paid \$25) tended to feel that the 10-hour subject should be paid \$30.29. In contrast, those who were told that they had worked 10 hours felt they should be paid \$35.24. The source of the bias can be seen in the bimodality of the distribution of perceived "fair" wages. One mode was at equal pay (\$25 for both) while the other mode was at equal hourly wage (so the 10 hour worker gets paid approximately \$35.70). Interestingly, the difference between the two treatments was mainly in the proportion at each mode. Those who had worked 7 hours showed more subjects at the equal pay level mode while those who had been told they'd worked 10 hours showed more subjects at the equal hourly pay mode. In other words, both groups recognized two compelling norms: equal pay for equal work and equal pay for equal output. Yet their roles determined (in part) which norm they picked. Such conflicts could easily arise even if there's disagreement about measuring input levels (which often are not fully observed). More broadly, when there is not universal agreement about what is fair division, individuals trying to act "fair" may produce even more conflict than individuals acting in a self-interested manner.

Let us return to the case of teacher absenteeism. The PROBE report surveyed teachers extensively in many areas of India and noted high absenteeism levels. Its in-depth interviews are illuminating about their attitudes and highlights how teachers often feel unmotivated. Some of this discouragement can be viewed as a perceived failure of

reciprocity. As noted earlier, individuals strongly adhere to the norm of reciprocity. Failures of reciprocity (or perceived failures) can result in punitive or self-interested behavior in response. Teachers may feel a strong social preference early on and be motivated to teach and give much more than they need to. After all, from a pure self-interest motive, they know they can get away with very little teaching. Yet they may be initially motivated to do more, to come to school, to struggle with tougher students and so on. They may view these contributions as a “gift” in large part due to the initial framing of the job (as a “plum job, with good salaries, secure employment and plenty of other time for other activities”). Thus, a young teacher may think, “I am giving a lot to the school”. As with any giving, however, the teacher may expect strong reciprocity and see (perhaps in a self-interested way) many outcomes as a lack of reciprocity. For example, the Probe report notes that many schools have terrible infrastructure; accordingly, teachers may feel that the government is not reciprocating their “gifts”. This may be especially exaggerated by the transfer system in India which moves teachers to various areas, disrupting the lives of teachers. Thus both the benign neglect of schooling and the active transfers could easily drive teachers to feel that the government does not reciprocate their efforts. They may also come to feel similarly vis a vis parents, who they may feel do not care about their children’s education.

Even an initially motivated teacher may very quickly feel justified in their growing apathy. They gave it their best and think that their efforts were not reciprocated. Are these inferences justified? Perhaps not. As in the Messick and Sentis study teachers may very well be making such inferences in a self-interested way. The failure of the context may very well be in it allowing teachers to make such biased attributions of fairness. Alternatively, teachers may very well be justified in these attributions. We simply cannot tell.

At a deeper level, these studies of fairness suggest that the problem of corruption may have interesting social preference wrinkles. People may be more willing to avoid taxes if they feel they are not “fair”. This judgment of unfairness could be the result of getting very few government services or having to bribe corrupt middleman in order to procure

government services. Economic models of corruption, by assuming blatant self-interest, ignore the tension corruption generates. If most people, as the evidence suggests, have strong social preferences, then corrupt acts will require self-justification. This resembles anecdotal evidence on corruption experiences. It is very rare that an official simply asks for a bribe. The request is often couched with an explanation for the reason for the bribe. Even though the bribe is clearly a violation of the law, there is usually a story that serves to justify it in the context of the law. For example, a customs officer may point to improper packaging as a reason for an extra payment. These kinds of insights may one day help to better understand the nature of corruption and anti-corruption policies.

Psychology of the Poor

The work I have described so far starts with a very simple premise: apply the general insights of psychology to the context of development. This raises the question of whether there are psychological features that are unique to poverty. Are there any judgmental heuristics, for example, that are uniquely applied in some developing countries? Are there cultural differences in some developing countries that may in turn reflect in their psychology and hence on economic outcomes? These questions all fall under the purview of cultural psychology, a long dormant branch of psychology that is now enjoying resurgence (Matsumoto 2001). For example, recent work suggests that East Asians are more overconfident than European-Americans in making probabilistic estimates and tend to focus less on individuals and more on contexts in making attributions (see Nisbett 2004). Though the formal explanations of such cultural differences are topics of ongoing discussion, lab experiments are identifying a seemingly culture dependent difference in reasoning. Given the large differences in context and affective experiences between rich and poor, might there be equally large differences between these two “cultural” groups as well? There is currently no work that I know of but it seems a question worth pursuing.¹⁸

¹⁸ In on-going work with Marianne Bertrand and Eldar Shafir, we have isolated a few suggestive differences between the urban poor and the well-off in decision making. I should note here that this type of work should not be confused with pejorative work that argues that the poor are “dumber”. Instead, this work attempts to highlight more subtle differences in cognitive heuristics used by the poor and the well-off.

Concluding Observations

Much of recent development economics has stressed the importance of institutions. Property rights must be enforced to give appropriate incentives for investment. Government workers must be given appropriate incentives to ensure high quality public service provision. Banking may need to be privatized to ensure a well-functioning credit system that in turn allows for better savings and smoother consumption. The common theme here is that institutions must be improved to help to resolve issues between people. They may reduce externalities, solve asymmetries of information, or help resolve coordination problems. This focus on institutions resolving problems *between* people, rather than *within* a person is natural to economists. The predominant economic model of human behavior leaves little room for individuals themselves to make mistakes. In fact, economists assume that people are unbounded in their cognitive abilities, unbounded in their willpower and unbounded in their self-interest (Mullainathan and Thaler 2001). Once we admit human complexities, institutional design in development becomes not just about solving problems *between* people. It is also about developing institutions that help any one person deal with their own “problems”. I hope the small set of examples I have presented illustrate how a deeper understanding of the psychology of people might eventually improve development economics and policy.

References

- Akerlof, George A, 1991, “Procrastination and Obedience,” *American Economic Review*, 81 (2) pp. 1-19.
- Ainslie G., 1992. *Picoeconomics*. Cambridge: Cambridge University Press.
- Ardener, Shirley, and Sandra Burman, eds., 1995. *Money-Go-Rounds: The Importance of Rotating Savings and Credit Associations for Women*. Washington DC: Berg Publishers.
- Aronson, J., Steele, C. M., Salinas, M. F. Lustina, M. J., 1998, “The effects of stereotype threat on the standardized test performance of college students,” in E. Aronson, (Ed.), *Readings About the Social Animal* (8th edition). New York: Freeman.
- Ariely, D., and Klaus Wertenbroch, 2002, “Procrastination, Deadlines and Performance: Self-Control by Precommitment,” *Psychological Science*, 13, pp. 219-224.

Becker, G., 1993, *Human Capital*. 3rd ed. Chicago, IL: University of Chicago Press.

Asharf, N., Dean Karlan and Wesley Yin, 2004, "Tying Odysseus to the Mast: Evidence from a Commitment Savings Product," mimeo, Princeton University.
<http://www.wws.princeton.edu/~dkarlan/downloads/SEED.pdf>

Babcock, L., and Loewenstein, G., 1997, "Explaining bargaining impasse: the role of self-serving biases." *Journal of Economic Perspectives*, 11, 109-126.

Bertrand, M., Sendhil Mullainathan and Eldar Shafir, 2004, "Behavioral Economics of the Poor," *American Economic Review*, May.

Besley, T. and A. Case (1994). "Diffusion as a Learning Process: Evidence from HYV Cotton." Working Paper.

Camerer, C. and Roberto Weber, 2003, "Cultural conflict and merger failure: An experimental approach," *Management Science*, 49(4), 400-415.

Conley, Timothy and Christopher Udry, "Learning About a New Technology: Pineapple in Ghana", Mimeo, Yale.

Darley, J.M., and B. Latané, 1968, "Bystander intervention in emergencies: Diffusion of responsibility," *Journal of Personality and Social Psychology*, 8, pp. 377-383.

Duflo, Esther, Michael Kremer and Jonathan Robinson, 2004, "Understanding Technology Adoption: Fertilizer in Western Kenya, preliminary results from field experiments," mimeo, MIT and Harvard University.

Farkas, Steve and Jean Johnson, 1997. *Miles to Go: A Status Report on Americans' Plans for Retirement*. New York, NY: Public Agenda.

Foster, Andrew and Mark Rosenzweig, "Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture," *Journal of Political Economy*, Vol. 104, 1996, pp. 1176-1209.

Frederick, S., George Loewenstein, G. and Ted O'Donoghue, (2002), "Time Discounting and Time Preference A Critical Review," *Journal of Economic Literature*, 40(2), 351-401.

Genesove, David and Chris Mayer, 2001, "Loss Aversion and Seller Behavior: Evidence from the Housing Market," *Quarterly Journal of Economics* 116(4), 1233-1260.

Gilbert, D. T., R. W. Tatarodi and P. S. Malone, 1993, "You can't not believe everything you read," *Journal of Personality and Social Psychology*, 65, 221-233.

- Gruber, Jonathan and Mullainathan, Sendhil. "Do Cigarette Taxes Make Smokers Happier," mimeo, MIT, March 2002. <http://econ-www.mit.edu/faculty/gruberj/files/happy81.pdf>
- Grace, N. D., H. Muench, and T. C. Chalmers, 1966, "The Present Status of Shunts for Portal Hypertension in Cirrhosis," *Gastroenterology*, 50, pp. 684-691.
- Gugerty, Mary Kay, 2001, "You Can't Saving Alone: Testing Theories of Rotating Savings and Credit Organizations," mimeo, Harvard University.
- Güth, Werner, Rolf Schmittberger, and Bernd Schwarze (1982) "An Experimental Analysis of Ultimatum Bargaining," *Journal of Economic Behavior and Organization*, 3:4 (December), 367-388.
- Hastorf, A. and H. Cantril, "They saw a game: A case study," *Journal of Abnormal and Social Psychology*, 49, pp. 129-134.
- Heath, C.; Bell, C.; and Sternberg, E. "Emotional Selection in Memes: The Case of Urban Legends," *Journal of Personality and Social Psychology* 81, no. 6 (2001)
- Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., Gintis, H., & McElreath, R., 2002, "Cooperation, Reciprocity and Punishment: Experiments from 15 small-scale societies." Book manuscript.
- Hoff, K. and Priyanka Pandey, 2004, "Belief Systems and Durable Inequalities: An Experimental Test of the Effect of Indian Caste on Performance," World Bank Working Paper.
- Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler, 1990, "Experimental Tests of the Endowment Effect and the Coase Theorem," *Journal of Political Economy*, 98, pp.1325-1348.
- Jacobs, R.C, and D.T. Campbell, 1961, "The perpetuation of an arbitrary tradition through several generations of a laboratory microculture," *Journal of Abnormal and Social Psychology*, 34, pp. 385-393.
- Laibson D., 1997 Golden Eggs and Hyperbolic Discounting. *Quarterly Journal of Economics*. 62, 443-478.
- Madrian, B. C. and D. F. Shea, 2001, "The power of suggestion: Inertia in 401(k) participation and savings behavior," *Quarterly Journal of Economics*, 116(4), 1149-1187.
- Matsumoto, David, *The Handbook of Culture & Psychology*.
- Messick, D. M., & Sentis, K. P. (1979). "Fairness and preference," *Journal of Experimental Social Psychology*, 15, 418-434.

- Miguel, Edward and Michael Kremer, 2003, "Networks, Social Learning, and Technology Adoption: The Case of Deworming Drugs in Kenya", mimeo, UC-Berkeley and Harvard University.
- Montgomery, Mark and John Casterline, 1998, "Social Networks and the Diffusion of Fertility Control," Working Paper.
<http://www.popcouncil.org/pdfs/wp/119.pdf>
- Mullainathan, S. and Richard Thaler, 2001, "Behavioral Economics," *International Encyclopedia of Social Sciences*. Pergamon Press, 1st edition, October: 1094-1100.
- Munshi, Kaivan (2003), "Social Learning in a Heterogeneous Population: Technology Diffusion in the Indian Green Revolution", mimeo, Brown University.
- Odean, Terry, "Are Investors Reluctant to Realize Their Losses?" 1998, *Journal of Finance*, pp.1775-1798.
- O'Donoghue T, and Matthew Rabin, 2001, "Choice and Procrastination," *Quarterly Journal of Economics*. 66, 121-160.
- Phillips, David P. 1979. "Suicide, Motor Vehicle Fatalities, and the Mass Media: Evidence toward a Theory of Suggestion", *American Journal of Sociology*, 84, 1150-1174.
- Phillips, David P. 1980. "Airplane Accidents, Murder, and the Mass Media: Towards a Theory of Imitation and Suggestion", *Social Forces*, 58, 4, 1001-1024.
- Powers, E., and H. Whitmer. 1951. *An experiment in the prevention of delinquency*. New York: Columbia.
- Read, D., George Loewenstein, and S. Kalyanaraman, 1999, "Mixing virtue and vice: Combining the immediacy effect and the diversification heuristic," *Journal of Behavioral Decision Making*, 12, 257-273.
- Regan, R.T., 1971, "Effects of a favor and liking on compliance," *Journal of Experimental Social Psychology*, 7, 627-639.
- Retherford, Robert and James Palmore. 1983. "Diffusion processes affecting fertility regulation." In Rodolfo A. Bulatao and Ronald D. Lee (eds.), *Determinants of Fertility in Developing Countries*, vol. 2. New York: Academic Press.
- Rogers, Everett, 2003, *Diffusion of Innovations*, 5th edition. Free Press.
- Ross, L. and Richard E. Nisbett, 1991. *The Person and The Situation: Perspectives on Social Psychology*. McGraw Hill.
- Samuelson, W., and Richard Zeckhauser, 1988, "Status quo bias in decision making." *Journal of Risk & Uncertainty*, v. 1, pp. 7-59.
- Sherif, M., (1937), "An experimental approach to the study of attitudes," *Sociometry*, 1, pp.90-98.

Strotz, R. (1956). "Myopia and Inconsistency in Dynamic Utility Maximization." *Review of Economic Studies*, 23, pp. 165-180.

Thaler, R., "Some Empirical Evidence on Dynamic Inconsistency," *Economic Letters*, 8: 201-207.

Thaler R., 1988, "The Ultimatum Game," *Journal of Economic Perspectives*.

Thaler, R., and Shlomo Bernatzi, "Save More Tomorrow: Using Behavioral Economics to Increase Employee Saving," *Journal of Political Economy*, forthcoming.

Trope, Y., and A. Fishbach, "Counteractive self-control in overcoming temptation," *Journal of Personality and Social Psychology*, 79, pp. 493-506.

Vermeesch, Christel, 2003, "School Meals, Educational Achievement and School Competition: Evidence from a Randomized Evaluation," working paper, University of Oxford.
<http://hicks.nuff.ox.ac.uk/users/vermeesch/schoolmeals.pdf>

Wartenbroch, K., 1998, "Consumption self-control by rationing purchase quantities of virtue and vice," *Marketing Science*, 17, pp. 317-337.