

Datta, Saugato; Mullainathan, Sendhil

Working Paper

Behavioral design: A new approach to development policy

WIDER Working Paper, No. 2014/103

Provided in Cooperation with:

United Nations University (UNU), World Institute for Development Economics Research (WIDER)

Suggested Citation: Datta, Saugato; Mullainathan, Sendhil (2014) : Behavioral design: A new approach to development policy, WIDER Working Paper, No. 2014/103, ISBN 978-92-9230-824-7, The United Nations University World Institute for Development Economics Research (UNU-WIDER), Helsinki,
<https://doi.org/10.35188/UNU-WIDER/2014/824-7>

This Version is available at:

<https://hdl.handle.net/10419/103001>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



WIDER Working Paper 2014/103

Behavioral design

A new approach to development policy

Saugato Datta¹ and Sendhil Mullainathan²

August 2014

Abstract: Successful development programs rely on people to behave and choose in certain ways, and behavioral economics helps us understand why people behave and choose as they do. Approaching problems in development using behavioral economics thus leads to better diagnosis and to better-designed solutions. This paper sketches how to use behavioral insights to design development programs. It distills key insights of behavioral economics into a simple framework about the constraints under which people make decisions. It then shows how this framework leads to a set of behavioral design principles whose use can improve the reach and effectiveness of many development programs.

Keywords: behavioral economics, development, program design

JEL classification: D03, D04, O12, O13, O14

¹ideas42; ²Harvard University; corresponding author: Saugato@ideas42.org

This study has been prepared within the UNU-WIDER project on 'New Approaches to Measuring Poverty and Vulnerability', directed by Jukka Pirttilä and Markus Jäntti.

Copyright © UNU-WIDER 2014

ISSN 1798-7237 ISBN 978-92-9230-824-7

Typescript prepared by Anne Ruohonen at UNU-WIDER.

UNU-WIDER gratefully acknowledges the financial contributions to the research programme from the governments of Denmark, Finland, Sweden, and the United Kingdom.

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.

UNU-WIDER, Katajanokanlaituri 6 B, 00160 Helsinki, Finland, wider.unu.edu

The views expressed in this publication are those of the author(s). Publication does not imply endorsement by the Institute or the United Nations University, nor by the programme/project sponsors, of any of the views expressed.

1 Introduction

1.1 Behavioral economics and development

Since the 1970s, the discipline of behavioral economics has brought together psychologists, economists, brain scientists and others in a quest to understand human behavior better. Its synthesis of insights from each of these disciplines has clarified our understanding of phenomena that are hard to explain satisfactorily using the tools of any one.

As a result, behavioral economics has changed the way we think about why people choose as they do and what motivates their decisions and actions. This has complemented valuable insights from neoclassical economics, particularly those arising advances in our understanding of asymmetries of information and mechanism design, which have helped us understand why markets and governments so often fail and what can be done to remedy such failures. Behavioral insights should not be thought of as overturning these insights. In some cases, they do conflict with them. But in many cases, they are complementary, helping further our understanding and expanding our policy choices.

Behavioral economics has provided new answers and new approaches to important questions in many areas of economics. Within development, for example, it has provided fresh insights into crucial questions about why the poor stay poor (see Banerjee and Mullainathan 2008; Bernheim et al. 1999, 2013). It has furthered our understanding of the nature of poverty by showing how poverty is as much about psychological and cognitive scarcity as about financial and material deprivation, which are the focus of traditional economic theorizing about poverty (see Bertrand et al. 2004). Both in this instance and in others, the insights provided by behavioral economics alert us to the importance of factors and mechanisms beyond those—such as the critical importance of access to information or the effects of uncertainty on decision-making—highlighted by neoclassical economics. In so doing, they have expanded our understanding both of important policy problems and of the tools available to tackle them.

As in the case of neoclassical economics, whose insights have revolutionized everything from how governments handle information to how teachers and nurses in developing countries are paid, these behavioral insights have not stayed theoretical. They have been used to design innovative solutions to persistent problems in development, especially those (such as of uptake, adoption, utilization, regularity of use, etc.) that remain unresolved after problems of provision, access and pricing have been tackled using the tools of traditional economic policy, such as taxes, subsidies, the provision of information, etc. Behavioral economics' understanding of such behaviors has informed the design of financial and health products for the poor in the developing world (see Ashraf et al. 2006; Brune et al. 2011) and is increasingly being used to find novel ways to solve problems ranging from inadequate drug adherence to the slow adoption of new technology in farming or industry (see Duflo et al. 2011; Hanna et al. 2012), and could in principle be applied to many other areas.

1.2 Behavioral design for development: a roadmap

This paper focuses on these existing and potential applications of behavioral economics to program design, which we argue are key components of an emerging science of design for development. We develop this argument in the next section, where we argue that such a science of design complements recent advances in the rigorous evaluation of development programs. This 'evaluation revolution' has made it possible to measure *whether* a given program or policy works,

thereby turning the spotlight on the question of *how to go about designing programs that are likely to work*. Behavioral economics, we argue, is a key part of the answer to this question.

Such a ‘behavioral revolution’, we argue in Section 2.2, is already in play in policy in the developed world. But behavioral problems very similar to those this revolution addresses are also at the heart of many policy challenges in developing countries. In Section 2.3, we discuss how this provides the rationale for thinking about the potential role of behavioral insights in helping to shape development programs and influence development policy.

The goal of the remainder of this paper is then to show *how* behavioral economics can help design more effective programs and interventions to tackle persistent problems in many areas of development. We begin with an example. In section 3, we see how approaching a specific problem in agricultural policy through a neoclassical lens helps us work through some of what may be going on, but nevertheless leaves several questions unanswered. We then show how using a behavioral lens enhances our understanding of the problem and leads to ideas for fresh solutions.

In Section 4, we develop a parsimonious way to place the contribution of behavioral economics in perspective, arguing that its key insights can be thought of as allowing us to identify an expanded set of scarce resources that drive human behavior, beyond the financial or physical scarcities that are emphasized in neoclassical economics. We delve into the rationale for emphasizing each of these scarce mental resources, and show how doing so provides key insights into policy-relevant questions in development, ranging from drug adherence to labor productivity to the adoption of technology. In each case, we see how taking a behavioral view expands both our understanding of the problem and the reasons why programs sometimes fail, while also providing tools to tackle the problem at hand beyond those suggested by neoclassical economics.

In the final section, we therefore sketch some key elements of the science of behavioral design. Specifically, we discuss the importance of behavioral diagnosis, i.e. the identification of the key psychological bottleneck leading to a particular outcome (Section 5.1), and then condense the lessons from the literature into a set of principles that can be employed to address some common bottlenecks that arise in a variety of contexts (Section 5.2).

Of course, the proof of the pudding lies in the eating. Much of the evidence we present here comes from small-scale research projects designed specifically to test the validity or relevance of a specific behavioral insight to a particular context. The next step—and an important one—lies in learning about how these insights and the innovations resulting from them can be applied at scale in the real world. This will require behavioral innovations to be deployed in the real world, and their adoption and use followed.

We therefore end with a description of how we think the insights of the research described in this paper can be put to work to solve problems at scale. This can happen through modifications to the design of public programs or those run by non-profit organizations. In some cases, where there is scope for a private provider to break even by solving a behavioral problem, there may be scope for incorporating behavioral design into products that are sold on the market or offered by for-profit providers. But, as we emphasize in Section 6, this will require a new approach to identifying behavioral bottlenecks to program success or product adoption and a willingness to embed behavioral science in design from the ground up.

Finally, it is important to reiterate what this paper is *not* about. We do not claim that the insights and tools we discuss are applicable exclusively to problems in development. Rather, they arise from features of human psychology that are as applicable to the affluent as to the poor. Indeed, this is why they can be (and have been) used to solve similar problems in developed countries, as we

discussed while motivating the issue of their application to development. This is also why insights from policy experiments in developed-country settings can be helpful in thinking about the problems we tackle here.

Related, none of the research we draw upon, or any of the insights that we glean from it, should be read as marking out ways in which the poor in the developing world are inherently ‘different’ from their counterparts elsewhere. Rather, it is worth emphasizing that while the people affected by the kinds of programs we discuss in this paper live under a very different set of *economic* constraints than their counterparts in richer countries, they face very *similar cognitive and psychological* constraints as their richer counterparts—but with fewer institutional aids to help them overcome the consequences of these common constraints. Behavioral design, we believe, can help fill some of these institutional gaps.

2 Designing better development programs: the role of behavioral design

2.1 The evaluation revolution

The last two decades have been exciting ones for development policy. Scientific advances in evaluation—often equated with randomized control trials—have sparked enthusiasm and optimism about tackling the persistent problems of poverty. Policymakers now feel better equipped to judge whether their policies work. In theory—and increasingly, in practice—funding can be based on solid evidence. All is not settled, of course. Debates continue on how these evaluations fit into the bigger picture. But few would dispute that the conversation has changed. ‘Does it work?’ is now a question everyone asks. This question leads to a new challenge. Evaluations make it clear that some interventions are effective, but that many others are not. This leads to a new question: ‘How do we design development programs so that they work?’ The answer to this question lies in the emerging science of design for development.

This science has a key role for behavioral economics. Many programs, policies or interventions stumble primarily because of the way people behave. Programs are taken up less enthusiastically than those designing them expect, resources are not spent as anticipated, programs are not implemented as those devising them had intended or assumed, and so on. In such cases, behavioral economics provides useful explanations¹ that can form the basis of useful policy interventions. It helps us understand why these behavioral challenges occur. Better understanding leads to better diagnosis, which in turn leads to better-designed solutions.

2.2 The behavioral revolution in policy

To see how, consider first the ways in which behavioral economics is already showing up in policy in the developed world. One of the most prominent examples involves advances made in boosting retirement savings in the United States. The behavioral insight that the status quo has a big effect on behavior has led to a critical design change: instead of having employees check a box to enroll in a savings plan, employers can have them check a box to *not* enroll. Research shows that this

¹ Of course, neoclassical economics also provides several categories of explanations for such phenomena. For instance, take-up of a program or product could be low for a number of reasons: information may be absent, the real costs of participation may be higher than apparent to an outside observer or researcher or, conversely, the benefits may be lower (or thought to be lower) than anticipated. Yet as we will show below, there are many situations where these or similar explanations sometimes take us only part of the way.

small change can boost savings by over 40 per cent. US pension policy has changed in response. Firms can now default people into savings and many do.

In the United Kingdom, a ‘Behavioral Insights Team’ has experimented with and discovered cost-effective ways to increase tax compliance and reduce enforcement costs in the judicial system (UK Cabinet Office 2012). Behavioral economics has affected the design of regulatory policy in several countries, because it changes how we think about people’s understanding of complex contracts (such as cellular phone plans) or how they process complicated information (such as the implications of their credit card bill). The European Union (see Ciriolo 2011) has used this understanding to alter competition policy. In the United States this has led to changes in the way credit cards bills must describe the balance due—no longer reporting only the minimum but also how long it would take to pay off the balance at that amount (often the answer is ‘forever’).

All this is possible because behavioral economics drastically changes how we think about people. It helps us realize that the assumptions we make—sometimes without realizing—when we design programs do not match the way people actually make decisions. Our intuitions—and those in economic models—overlook many of the important things that make people tick. Some of the insights seem self-evident: for example, that we have self-control problems. Yet even these seemingly self-evident insights are often overlooked: we may all know we have self-control problems, but we forget this fact when we design programs. This applies equally to development programs, to which we now turn our attention.

2.3 Behavioral economics and development programs

Successful development programs—like other kinds of programs—rely on people to behave and choose in certain ways. For schools to be effective, parents must ensure their children show up. For vaccination programs to work, parents must bring their children to vaccination stations. For new inventions—whether insecticide-treated bed-nets, smokeless indoor stoves, or new seeds—to provide benefits, people must adopt and use them. Similarly, people’s health depends on the health care available to them. But for this care to be effective they must also do certain things. They must make prenatal visits. They must adhere to drug regimens. They must feed their children effectively. They may need to clean the water they drink.

Behavior thus affects whether the provision of schools, healthcare, improved seeds, or other technologies have the effect they are intended to have. Because behavior matters to much, programs work better when they are designed to match people’s actual psychology, with due consideration to the way they make decisions. Even programs based on sound medicine, agronomy or education science can fail if their design makes the wrong assumptions about behavior. There is thus enormous scope to apply behavioral design to development programs.

This paper outlines how to use behavioral insights to design better development policy. Doing so requires first answering one question. Is there anything special about the psychology of the poor that makes them different in some way? Recent research suggests an important way in which there is. The conditions of poverty themselves create additional psychological burdens (see Mani et al. 2013; Shah et al. 2012). In a few years, we anticipate (hope) these psychological insights will translate into policy. For now, though, we focus on the powerful insights, and designs come from focusing on the psychological phenomena that affect all of us, whether rich or poor. Of course the consequences of these psychologies may be different for the poor than the better off (see for example Bertrand et al. 2004). But the underlying mechanisms can be thought of as roughly the same. This pragmatic approach may be incomplete, but it points us to many low-hanging policy fruit, where a behavioral-economics lens can improve our understanding of when and why

programs work well or poorly, and suggest new ways to tackle situations where there is scope for improvement.

Behavioral economics affects program design in three steps. First, it changes how we *diagnose* problems. For example, when we see parents failing to vaccinate their child we may be tempted to conclude that they do not understand the value of vaccination. Behavioral economics forces us to consider another possibility: they want to vaccinate, they understand the benefits, but they don't get around to doing it. Vaccination may be one of many behaviors, such as savings or going to the gym, where what we do fails to match up with what we want to do. Second, it changes how we *design solutions* to problems. In some cases it may suggest that something as simple as a reminder can have an unreasonable impact on behavior. In others it may suggest a different way to offset our tendency to plan our spending poorly. Finally, it changes how we *define* the scope of the problem. Problems we overlooked may suddenly become interesting ones to solve. We often focus on access ('Make sure people get the drugs they need at low cost'). Behavioral economics points us towards important problems that remain even after access is solved ('Make sure that people actually take the drugs they are given'), and provides ideas about how to tackle them.

3 An example: the behavioral economics of fertilizer use

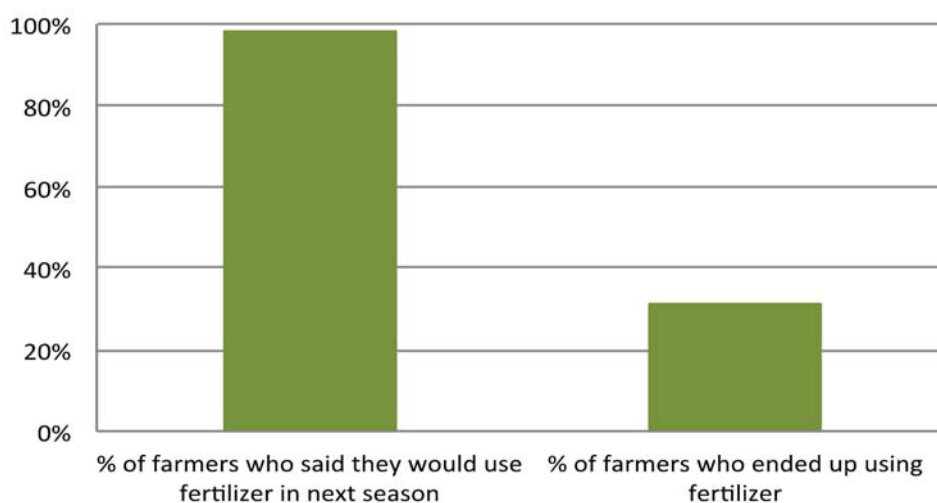
Working through a development problem will illustrate each of these steps. Many farmers in Sub-Saharan Africa use little or no fertilizer on their fields. This might explain why African crop yields lag behind those in Asia, where fertilizer use is higher (Morris et al. 2007). But why is fertilizer use in Africa so low?

3.1 The usual explanations

How do we diagnose this problem? Neoclassical economics offers several explanations, each of them plausible. One possible reason for low fertilizer use is that fertilizer is not easily available. Another is that fertilizer is too expensive. It is also possible that fertilizer does not work as well on real fields in real conditions as it does on test farms, so that it is the policymaker who is mistaken in recommending higher doses than are actually used. Finally, it is possible that fertilizer does increase yield, but that farmers do not know about this.

Each of these diagnoses is based on a presumption. The economic approach leads us to assume that a farmer who does not use fertilizer did not want to use any. We impute intentions from actions. When someone fails to do something we assume they were unwilling or incapable of doing it. So we try to understand why they cannot (fertilizer is not available) or do not want to (they do not understand the returns, or the policymakers' estimate of returns is flawed). These questions in turn affect the solutions we try. We focus on increasing access. We try to make fertilizer cheaper by subsidizing it. We try to inform farmers about the benefits of fertilizer to make them change their minds about using it. Or we conclude that there is in fact no sub-optimality: farmers are in fact doing what is right; it is the policymaker or agronomist who is mistaken.

Figure 1: Drop-off between intention and action



Source: Based on Duflo et al. (2010).

3.2 The behavioral approach: a different diagnosis

These are not bad questions to ask, of course. There are situations in which one or other of these diagnoses is the relevant one. The mistake is the other questions we fail to ask.

In the case of fertilizer, the questions we do ask may not lead far. Of course, there are places where fertilizer is not available, but there is considerable evidence that fertilizer fails to be used even when it is readily available, including in parts of sub-Saharan Africa where use is low but availability is ample. Second, fertilizer is sold in small amounts that even small farmers can afford (see Duflo et al. 2011) so that price per se is not the key bottleneck. Experiments show that fertilizer increases output not just on test farms in agricultural extension centers but under real conditions on real fields and without farmers making any other changes to their practices: for example, it raises net incomes of maize farmers in Western Kenya by as much as 18 per cent per season (see Duflo et al. 2010). And even farmers who do not use fertilizer are aware of its benefits (Duflo et al. 2011). So fertilizer is available, affordable, effective, and appreciated. But it is still not used.

Behavioral economics leads us to a very different question. Our intentions do not always translate into actions. Sometimes we want to do things but do not do them. Many of us wake up later and go less often to the gym than we would like to. What if farmers have the same problem with fertilizer? In fact, the farmers themselves agree with this sentiment. About 97 per cent of Kenyan farmers surveyed by Duflo et al. (2011) said they intended to use fertilizer on their fields the following season, but only 37 per cent actually ended up using fertilizer. Asking the kinds of diagnostic questions economists *do not* usually ask leads to some interesting data. These data do not prove the case. And we must tread carefully with such casual questions; after all people may simply be telling us what we want to hear. But what is most important here is that a new possibility, a new diagnosis, is added to the mix. One that has some validity, at least on the face of it: some reason to believe it *might* be right.

There is of course nothing special about the two-thirds of Kenyan farmers who say they will use fertilizer and then do not. We all repeatedly fail to live up to our own intentions in big ways and small. It is merely the case that Kenyan farmers do too. The question is why.


Behavioral economics provides several possible answers. First, we all tend to think we will be much more willing to do things in the future than we actually are. So we *procrastinate repeatedly* about doing even things we want to do. We put things off to a tomorrow that never comes. In the case of Kenyan farmers, going to the market to get fertilizer is a bit painful, both in terms of time and money. Farmers plan to ‘do it tomorrow’, except that the planting season arrives before ‘tomorrow’ does (Duflo et al. 2011).

Second, we all lack self-control. We succumb to immediate temptations. Tomorrow we plan to cut back on sweets. When tomorrow becomes today we eat dessert. This is as true of farmers in Africa as it is of us. When the farmer is rich with cash at harvest they want to spend the money on fertilizer. By the time planting time arrives the money is no longer there (see Brune et al. 2011). The time between harvest and planting is a time when fertilizer can be overlooked. Money is spent on other things and farmers find themselves caught short later, when they need the money for farm inputs.

3.3 Behavioral solutions

A new diagnosis can lead to new solutions. If it is the small hassle cost of travel to town that leads to farmers procrastinating in buying fertilizer then home delivery early in the season should help raise use. Duflo et al. (2011) test this idea and find home delivery raises fertilizer use by 70 per cent. Think of how interesting and surprising this is: the alleviation of a small cost can have a large effect. In this case it is because the small cost was the snag that was causing procrastination.

Figure 2: Commitment accounts in Malawi



Note: Farmers who received commitment accounts purchased more inputs than those in control.

Source: Based on Brune et al. (2011).

Similarly, if farmers have self-control problems, we can give them a way to tie their own hands. They already *want* to save for fertilizer at harvest. We merely need a way to help them translate this intention into action: some way to keep the money ‘safe’ (from themselves) during the long period between harvest and planting. Suppose they had a special account that let them lock up some of their money, and free it for use at a time of their own choosing. This simple product works. Farmers given this option bought and used much more fertilizer and other inputs, leading to higher crop sales (Brune et al. 2011). We are not locking up their money. They are choosing to have the

option to do so². It was not information or fertilizer that was lacking. But there was no financial product that let them follow through on their desire to use fertilizer.

The fertilizer example shows a few lessons. First, we can often be blind to some diagnoses. Some snags simply seem too small: the cost of travelling to a nearby town, or the hassle of remembering to put some money aside soon after the harvest. Yet in some cases—as here—these overlooked diagnoses can be a big part of the story. Overlook them and we forgo a potentially powerful solution.

Second, the interventions are so powerful precisely because the snags are ‘small’. A large increase in income from using fertilizer can come from quite unexpected places. For example, early home delivery amounts to a 10 per cent discount on the market price of fertilizer, but it increases use by as much (70 per cent) as a 50 per cent subsidy would. Finally, these examples show the breadth of psychology. It is not merely about marketing or better tools of persuasion (sometimes it is that too). It is a deeper perspective on what makes people behave. Sometimes we can change behavior without ever changing people’s minds. In this case it is because they (or many of them) were already convinced.

4 The other limited resource

How do we incorporate the insights of behavioral economics to other problems beyond fertilizer use? Procrastination and self-control are just some potential psychological phenomena. What about the others? How do we do diagnose more systematically?

To help navigate the large set of findings, we condense the behavioral literature using one simple perspective about the *constraints* under which people make decisions. Economists and policymakers—indeed all of us—understand constraints all too well. Resources are limited: there is only so much money, time, staff, or even enthusiasm to go around. Yet we often do not realize that *mental* resources are also limited (Mullainathan and Thaler 2000). While we understand that physical resources must be carefully doled out, we are often blind to our finite mental resources. Without realizing, we often design programs assuming that people have unbounded cognitive capacity. We assume that they can think through complex problems effortlessly and quickly arrive at the ‘correct’ choice. We often assume unbounded self-control, which leads us to expect people to always resist temptations and do what they intend to do. These assumptions are often unstated, implicit, or even unconscious, but they show up in the way we design programs and policies.

Behavioral economics can be understood as identifying a few more limited resources. In practice, we have found it helpful to think about the limits on four basic mental resources, each of which we discuss below. For each, we will look at one problem. Then we will show a few other problems where thinking about each limited resource leads to new ideas about why problems occur and how we can solve them.

² In this case, the farmers benefited even without depositing much into the accounts. The very availability of the commitment product let them credibly claim that they had no ready cash when others (friends, neighbors, extended family members) wanted access to some of their harvest earnings.

4.1 Scarcity of self-control

4.1.1 Labor productivity in developing countries

Much of development depends on labor productivity. People's productivity affects their income, a village's overall level of output and firms' profits. Many programs focus on improving productivity. Though these programs are diverse they are based on a common set of diagnoses. Some programs focus on improving skills. They diagnose low productivity as a capacity issue: 'people are not capable of working more effectively'. Others focus on incentives, for example motivating teachers to show up at school. These diagnose low productivity as one of motivation: 'people have no interest in working harder'.

These diagnoses miss another possibility. Clearly, workers frequently work less hard than their employers would like. But is it possible that they also work less hard than they themselves want to, because they have self-control problems like the rest of us? Self-control is very hard, as anyone knows who has tried to stick to a diet in the face of a tempting but 'forbidden' dish. In the classic 'marshmallow experiments', children left alone in a room with a single marshmallow struggle visibly as they try to resist eating it in order to win a second marshmallow as a reward for their forbearance. Exerting self-control is physiologically effortful, leading to a faster pulse and decreased skin conductance (see Kahneman 2011). We now know that it makes sense to think of self-control as a psychic 'commodity' of which we have a limited stock (see Vohs et al. 2008), so that using some up for one task ('continuing to exercise when you want to stop') depletes the amount available for other tasks ('resisting the extra cookie after your workout').

Once we are attuned to the difficulty of self-control, we see work in a different way: as a series of tests of self-control. It takes self-control to identify, plan and execute all the tasks that need to be done, all the while resisting the many temptations and distractions that surround us. All of us sometimes lose some of these battles of self-control, resulting in our working *less hard than we ourselves would like to*. Kaur et al. (2011) thought it plausible that this was happening in an Indian data-entry company. They designed and offered workers a 'negative bonus' scheme. Under this scheme, workers were paid their usual piece rate if they met *self-chosen data entry targets* but penalized if they did not. About 35 per cent of workers chose to set non-zero targets, suggesting that they wanted to get themselves to work harder. These 'negative bonus contracts' increased output by an amount that was equivalent to the effect of increasing their piece rate by 33 per cent, and by *more than a year's worth of education*. Once again, a behavioral diagnosis led to a large impact: so large that it would have required raising wages by one-third. And the increase came from an unexpected source: not paying workers more, or training them afresh, but simply by giving workers a way to work as hard as they wanted to.

This insight could help solve other problems too. Policies to tackle high levels of absenteeism among public service providers (such as nurses and teachers) in developing countries usually rely on enforcement and monitoring. But if self-control problems among such workers prevent them from working as hard or showing up as often as they themselves would like, finding ways to mitigate self-control problems may lead to effective alternative solutions.

4.1.2 Self-control and problems in farming and saving

When we think about how to increase output on farms, we usually think about how to increase the adoption and use of inputs like fertilizer or improved seeds. These are clearly important. But there are other important behaviors that affect how productive farms are. For example, farmers weed much less than they should. For example, Banik et al. (2006) find that weeding twice a season

instead of once raises yields in India by 23 per cent for wheat and an enormous 49 per cent for chickpeas.

But why do farmers weed so little? We usually diagnose the problem as arising from a lack of knowledge: farmers may just not know how important weeding is. But weeding is also time-consuming, easy to postpone, and tedious. In sum, doing it requires self-control. So do other things farmers do not do enough of, such as channeling runoff water correctly. So behavioral economics suggests a different diagnosis, one of self-control problems. Once we recognize the role of self-control in agricultural work, we see many ways to increase productivity on farms beyond increasing input use.

Similarly, when we see low savings rates, we usually assume that people do not want to save, or that savings programs are not lucrative enough. As a result, we try to raise awareness about the importance of saving or to make saving more financially rewarding, whether by increasing interest rates or matching contributions. But building up savings is a process with many steps, several of which demand self-control. It requires self-control to not spend money when it is available, but also to make and stick to plans to go to the bank and deposit it. Recognizing the role of self-control in making it hard for people to save as much as they want led behavioral economists to design a special savings account for clients of a bank in the Philippines. These accounts allowed them lock up funds in their own accounts until a self-specified goal had been reached. Nearly 30 per cent of the clients who were offered such a restrictive account opened one, and the effect on the savings balance after one year was an 81 per cent increase (Ashraf et al. 2006). Achieving an increase of this magnitude would have required an unfeasibly large increase in the interest rate offered. Once more, we see how small behavioral interventions can have almost unreasonably large effects.

4.2 Scarcity of attention

4.2.1 *Technology adoption in developing countries*

Adopting and becoming proficient at using new machines, inputs or techniques of production is a big part of the development process. Yet technology adoption can be frustratingly slow. Many beneficial techniques are not adopted at all, and workers in developing countries often use newer machines and techniques less efficiently than possible long after they get access to them. To speed up the process, governments often use programs (such as agricultural extension programs) intended to teach potential users about the benefits of new technologies and how to use them correctly. These programs vary in their details, but they all rest on a common diagnosis: a lack of knowledge about how to use technology. Yet mastering a new technology often requires more than just knowing about it or even a superficial understanding of what it involves. Rather, it requires a user to be especially *attentive* to some particular features or aspects of the technology. If he does not pay attention to the right things, a person is unlikely to become adept at using the new technology.

This would not matter if people noticed everything, as they would if they had unlimited attention. But psychologists have found attention to be a limited resource, just as self-control is. In one famous experiment that has been repeated many times, over half of those tested fail to notice a gorilla walk across a basketball court because they are paying attention (as instructed) to the number of passes made by the teams on court. Similarly, subjects in dichotic listening tests concentrate on one of two distinct audio streams—each emanating from one headphone—and remember very little about even the stream they are asked to pay attention to and essentially nothing about the other. Attention—both visual and aural—is even more limited than we usually realize. This leads to a different diagnosis. Because of limited attention, people are unlikely to notice all aspects of a new technology. And they may fail to notice precisely those dimensions of

a new technology or technique that actually matter, because these dimensions are ones they are used to ignoring.

A recent study of seaweed farming in Indonesia (see Hanna et al. 2012) shows how limited attention can slow down or prevent technology adoption. Seaweed farming is done using the ‘bottom method’ where the farmers attach ‘lines’ through wooden stakes driven into the shallow seabed near the shore. Raw seaweed from the previous harvest is then taken, cut into pods, and attached to these lines. Farmers tend to these pods when the tide is low, and harvest the seaweed after 30-45 days. A number of things could affect yield. These include the size of the seaweed pods, the distance between lines, and the distance between pods on a line. Maximizing production and income involves figuring out the optimal combination of these dimensions.

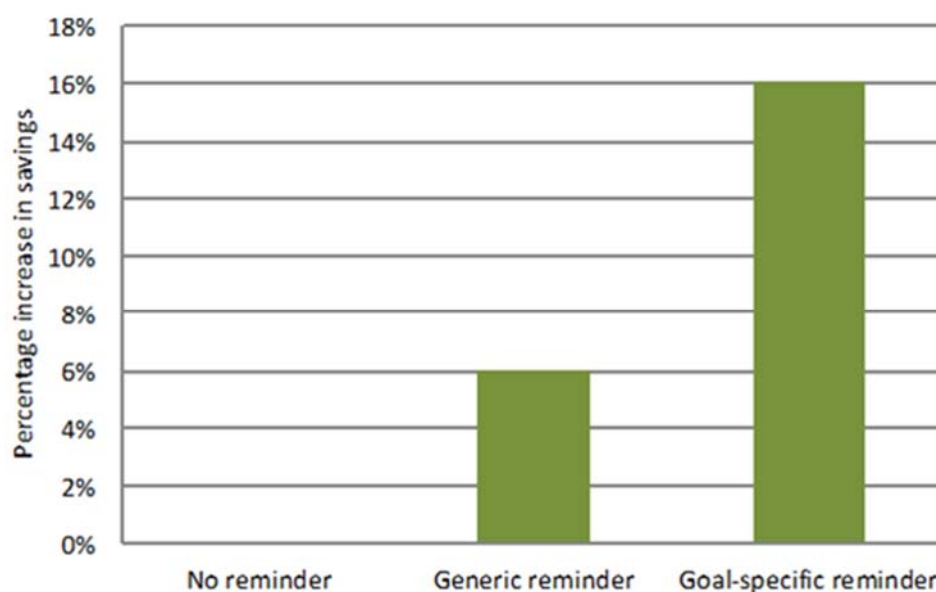
It turns out that farmers simply ignore pod size, which experiments have found does matter. When asked about the length of a typical line, the distance between lines, or the optimal distance between knots and lines, close to 100 per cent of farmers can answer the question. But when asked about the current pod size or the optimal pod size, only 16-17 per cent of them are able to provide an answer. Because seaweed farmers do not think pod size is important, they pay little attention to it. As a result, they also do not learn over time that it matters: it simply evades notice.

This diagnosis helps explain why simply providing information does not help. Even participating in experiments designed to measure the effect of pod size on output has no effect on farmers adopting the right pod size, although the experiments themselves show that moving from the worst to the best pod size raises income from seaweed farming by 30 per cent, and total income by 20 per cent. Just like people missed the ‘gorilla’, farmers are effectively blind to features of a technology they initially believed did not matter.

4.2.2 Further applications: saving and drug adherence

Behavioral economists have shown that limited attention prevents people from acting on their own intentions to save. So limited attention—and not the usual diagnoses of the absence of the desire to save or low returns to saving—may help explain why people save so little. If so, drawing people’s attention to their own plans to save may help increase savings rates. In a series of experiments in Peru, Bolivia, and the Philippines, behavioral economists successfully raised the amount people saved simply by providing them with timely reminders about their own saving goals (see Karlan et al. 2010).

Figure 3: Text message reminders increase savings



Source: Based on Karlan et al. (2010).

In the United States, close to 100,000 people have their limbs amputated each year because of complications arising from diabetes. Yet diabetes is easily treated, and most of the people who lose their limbs have been prescribed medication for diabetes, but do not take it regularly. Similarly, access to anti-retrovirals has made HIV/AIDS a manageable condition even for patients in many developing countries. But here, too, many people do not take their pills regularly. Tuberculosis has long been treatable if patients complete a course of medications—but many do not. In all these cases, not doing something trivially small—taking one’s medicines—has almost unimaginably damaging outcomes—the loss of a limb or even death.

Because the consequences of not taking medicines are so big, we usually assume a big gap in knowledge or understanding must be responsible. As a result, we typically focus on educational or communication campaigns about the importance of taking one’s medicines. Yet behavioral economics suggests that we might be missing a plausible alternative diagnosis: inattention. What if people simply forget to take their pills day after day? Many illnesses—including diabetes and tuberculosis after a certain stage—are not symptomatic most of the time, so it is all too easy to forget to take medication. This insight seems implausible, because it identifies something trivial—forgetting—as leading to something huge—such as a loss of life. Yet solutions to the problem of forgetting have proven very effective at raising drug adherence. In the United States, GlowCaps - pill bottles that light up if not opened at the right time—have dramatically raised adherence. In South Africa, simple text message reminders have been used to increase adherence.

4.3 Scarcity of cognitive capacity

4.3.1 Ineffective financial literacy programs

Running a small business involves a lot of financial management, ranging from keeping business accounts to managing debt. Yet many of those who run small businesses, in both developing and developed countries, are not financially literate enough to handle these tasks, causing their businesses to suffer and making it difficult for them to grow. In response, agencies and governments have developed financial literacy training programs for small business owners. However, there is little evidence that financial literacy training makes a difference to how small

businesses are run and how their owners manage their finances. For example, Drexler et al. (2011) evaluate a standard financial literacy program in the Dominican Republic and find that it has no impact on how entrepreneurs manage their finances, and consequently none on how profitable their businesses are.

This low impact is typically diagnosed as arising either from a lack of interest and motivation on the part of those who attend financial literacy trainings or from not enough useful material being covered in training sessions. So we often see interventions that aim to increase the number of sessions people attend, or to make the material covered in financial literacy training more comprehensive. Yet even these interventions have proven unsuccessful.

Behavioral economics can help us understand why. When we try to make financial literacy curricula more comprehensive and rigorous, we are assuming that people can process large amounts of complex information quickly and effortlessly. Yet research shows that this is simply not true. The cognitive resources available to people at any moment are limited and can be depleted by being used for other activities. So increasing the cognitive demands of financial literacy programs may in fact be making them less likely to succeed. These solutions target the wrong psychology.

On the other hand, behavioral economists have found that bounds on cognitive and computational ability lead us to ‘economize’ on cognition while making decisions. Wherever possible, we use fast, intuitive thinking or rough rules of thumb. An alternative way to try and improve outcomes from financial literacy programs is therefore to build them around simple rules of thumb of the kind that people actually use. A test of such a ‘rules-of-thumb’-based financial education program in the Dominican Republic found large increases in the adoption and use of good financial practices (such as separating business and personal accounts) as well as a jump in profitability (Drexler et al. 2011). Building solutions around psychology makes them more likely to succeed.

4.3.2 Further applications: pension and cash transfer programs

Many governments in low- and middle-income countries are concerned with getting more people to participate in pension systems and save for their old age. Yet few participate and those who do join programs often contribute very little. Governments typically diagnose the problem as arising from a combination of a lack of interest in saving for the future and the savings programs themselves not meeting people’s needs. So they try to make saving programs more attractive. One way in which they do this is by providing people with more choices about how much they can contribute, where their money is invested, and so on.

Yet behavioral economists points out that more choice is cognitively taxing, and giving people too many choices may overwhelm them. As a result, even those who want to save may end up not doing so because they find it too hard to choose between the many plans and rates available. Providing more options can actually make it less likely that any of them is chosen.

Thinking about choices in this way leads to very different solutions, such as simplifying plans and reducing the number of dimensions that people have to compare, or picking an option that people are automatically defaulted into. Choice simplification can be remarkably powerful. A study of loan take-up in South Africa found that reducing the number of combinations of interest rates and loan tenures led to as much of an increase in uptake as reducing the interest rate charged by 2.3 percentage points (Bertrand et al. 2010).

Figure 4: Behavioral manipulations increase loan take-up

Behavioral Manipulation	Reduction in Interest Rate Required for Same Effect on Take-up
Reducing number of loan options from 4 to 1	2.3 percentage points
Using female picture for male applicants	3.2 percentage points

Source: Based on Bertrand et al. (2010).

We also often diagnose the problem of low uptake of programs to a lack of interest in the program among those targeted, and respond by providing more information. Yet if cognitive space is limited, a flood of information may simply overwhelm people and could actually *reduce* uptake. In addition, scarcity of cognitive resources can lead people to attach more weight than expected to features of the way that information is *presented*, such as whether something is couched as a gain or a loss. Once we know that cognitive biases are at work, however, we can design communication and advertising to remove these biases or to take advantage of beneficial ones. Behavioral economics leads us to pay attention to features of communication—how changes are framed, how many options are presented, and how complex choices are—that we might otherwise dismiss as unimportant.

4.4 Scarcity of understanding

4.4.1 Under-use of oral rehydration solution

Over half a million infants throughout the developing world—around 150,000 of them in India alone—die of diarrhea each year. Yet by some estimates, over 90 per cent of these deaths could be easily averted through the use of a balanced solution of salts known as oral rehydration solution (ORS). Why is ORS not used enough?

Most attempts to solve this problem diagnose it as arising either because people do not know about ORS and how it works, or because it is not readily available or affordable. We therefore try to make sure that ORS is cheap and easy to find by working on distribution and cost, while also using advertising and other kinds of awareness campaigns to inform mothers about its benefits.

However, ORS use remains low even where it is available and cheap. Behavioral economics helps us understand why. In making decisions—such as what to do when a child has diarrhea—we have to rely on an underlying theory, possibly an unstated or implicit one, about the disease and its control. This constitutes our ‘mental model’ of the world. Usually, we assume that these underlying theories are broadly correct. But behavioral economists argue that this understanding, too, is scarce; not all underlying causal relationships are correctly or accurately understood.

A child with diarrhea is constantly leaking fluids. Given this, a perfectly plausible mental model of the disease would imply that putting any more liquids into the child will only make it sicker; keeping the child ‘dry’ is better. Indeed, when poor women in India are asked whether the solution to a child with diarrhea is to increase or decrease its fluid intake, 35-50 per cent say that the answer is to decrease it. But with this mental model of the disease, it will never make sense to use ORS—and unless this model is somehow overturned, ORS will not be tried. This helps explain why many mothers do not use ORS despite its being cheap, readily available and well known. It also suggests

that any effective solution will have to tackle the flawed mental model at its root: without doing so, information or exhortation is unlikely to have much effect.

4.4.2 Further applications: imbalanced fertilizer use and schooling decisions

Rice farmers in some parts of India over-use nitrogenous fertilizer. Usually, we think that this is the result of poor pricing policies and a lack of awareness about the right ratio between various kinds of fertilizers. This has led governments to concentrate on disseminating information on the right way to use fertilizer through extension offices or information campaigns. However, behavioral economists argue that a flawed mental model is at work here. Farmers are used to extrapolating the likely yield of a crop by looking at the extent of its green leafy growth: ‘green = healthy’). Nitrogenous fertilizer is good for such leafy growth. In the case of many crops, (like spinach), this provides a good rule of thumb. But in the case of grains, too much leafy growth can detract from yield. Relying on this otherwise useful mental model can thus lead farmers to use too much nitrogenous fertilizer.


Similarly, some poor parents pull their children, or some of their children, out of school after only a few years of schooling. The standard diagnosis of this problem is that parents lack interest in schooling or that schooling is too expensive. As a result, many governments have adopted policies involving reducing or eliminating school fees or providing financial assistance to poor parents as ways of increasing schooling among the poor.

However, there is evidence that flawed mental models may play a role here too. Data from the developing world shows that each additional year of schooling adds roughly as much to earning power as the previous one. Yet parents in both Madagascar and Morocco strongly believe that primary school is about half as valuable as secondary school, which is about half as valuable again as high school (Banerjee and Duflo 2011). Many parents therefore think of schooling as essentially worthless unless they can afford to send their children all the way through high school. So they may allow their children to drop out with even less education than they can afford to pay for. Indeed, just telling parents in Madagascar about the average income gains from spending one more year in school for children from backgrounds like their own increased test scores for their children, and was particularly effective for children whose parents had earlier under-estimated the return to education (Banerjee and Duflo 2011).

5 Behavioral design

Unlocking the potential of behavioral economics requires a systematic approach to design. We break the design problem into four parts: problem definition, diagnosis, design, and testing. Here, we focus on the two central parts of this scheme: diagnosis—*where to intervene* and design—*how to intervene*.

Figure 5: The stages of the behavioral design process



Source: Authors compilation.

5.1 Finding the behavioral bottleneck: diagnosis

The first step to finding a solution is identifying the root of the problem. Where are things going astray? We may be tempted to identify bottlenecks intuitively, but our intuitions about psychology can be fallible. In addition, there are usually many psychological barriers or phenomena that could be at play in a given situation. Guessing and eliminating using a series of experiments is thus likely to be both unreliable and expensive. We instead need a systematic approach to identifying *candidate* bottlenecks. This process, which we call ‘behavioral mapping’, reveals ‘behavioral stress points’, each of which is a possible intervention point.

Behavioral mapping begins with a problem. For instance, it might be that farmers apply too little fertilizer to their fields. It then decomposes this behavior into a number of parts. Did the farmer ever intend to buy the fertilizer at some point? If he did not, the bottleneck is at the point of decision. If he did, then the bottleneck is one of action—of following through on a decision—or of belief revision—sticking to the original decision. Figuring out which of these is most relevant requires us to use further diagnostic tools. Did the farmer continue to believe for long periods that he would use fertilizer tomorrow, just not now? Did the farmer simply forget? Note that each of these diagnostic questions is derived from a psychology we described above: the first looks for self-control problems, whereas the second is linked to limited attention.

There is, however, no mechanical, one-to-one mapping between a psychology and bottlenecks. For example, scarcity of self-control could imply bottlenecks at several stages. A farmer lacking self-control may not save the money to buy fertilizer. Or he might not take the effort to figure out if fertilizer is right for his farm—so that he never gets to the point of trying to save to buy it. Behavioral mapping is thus a process that generates questions; these then lead to surveys—qualitative and quantitative. The data from these surveys then guide the next set of questions. A careful use of data and observation then allows us to arrive at a shortlist of the most important among the hypothesized bottlenecks mapping has helped shortlist. These bottlenecks eventually feed into designs, which are then tested.

5.2 Behavioral design principles

The design phase—deciding *how to intervene* in a particular situation—follows this diagnosis. The designs obviously depend on the hypothesized bottlenecks. Solving a problem of forgetting makes little sense if the person never intended to take the action. Conversely, since the bottleneck guides the design, it is entirely possible that an intervention that works well for one problem has no impact at all on a different one. For instance, the provision of information may be useful where lack of understanding is an issue, but may do little where the key bottleneck arises from self-control problems. Diagnosis is thus critical; design follows from it.

However, once a relevant behavioral bottleneck has been identified, the set of design principles we describe below can guide behavioral design. Of course, these principles need to be applied judiciously. But they give a clear sense of how behavioral insights can improve the design of policies and interventions once we have honed in on the right bottlenecks and the reasons for them

Principle 1: Facilitate self-control by employing commitment devices

As we have seen, we all have self-control problems. We may decide to do one thing but then find it challenging to follow through. We also often seek ways to stick to our own plans. Signing up for a gym membership is one way to encourage our future selves to go to the gym. So is the more aggressive (and more expensive) decision to sign up for a training session.

Interventions or suitably designed products can affect this imbalance between intention and action. Clocky is an alarm clock that makes it easier to stick to plans to wake up early. When you hit its snooze button, Clocky jumps off the nightstand and rolls into a corner. When it goes off again, you have no choice but to get up. Similarly, a non-profit venture called StickK allows people to stake their own money on things they are trying to get themselves to do, such as give up smoking. Their funds are only returned to them if they are certified as having met their goal by a third party. So far, nearly 150,000 ‘commitment contracts’ have been taken out on www.StickK.com.

Poor people in the developing world suffer from the same self-control problems as the rest of us. They also try to find workarounds that make it easier for them to stick to their intentions. For example, poor women in developing countries often ‘borrow to save’ by taking a loan from a microfinance institution (MFI) and then place that loan into a savings vehicle. The point of doing this is that the need to pay the MFI back provides the discipline needed to save regularly, which is otherwise not available. Borrowing to save thus results in savings while saving on your own may result in nothing.

Policy design can incorporate these insights more explicitly. We saw this for savings in the Philippines (Section 4.1). Similarly, farmers in Malawi used restrictive accounts to get around their own tendency to spend harvest proceeds too quickly, which resulted in a cash crunch and under-investment in the next agricultural season (see Section 2.4). In both cases, the saver *valued* restrictions on future behavior.

There are many further applications. For instance, poor people in developing countries often take high-interest loans for predictable expenses they could easily save up for, such as for home repairs, appliances, school fees or medical expenses related to childbirth. Low- to middle-income women in rural India take on large debts to pay for the costs associated with delivering a child in hospital. An ongoing experiment in rural Andhra Pradesh, India, allows pregnant women and their families to save regularly in a designated account. The saved amount cannot be withdrawn until they give birth, but forms the basis of a loan to cover childbirth-related expenses. Take-up of this product in the pilot area suggests that it is fulfilling a deeply felt but unmet need.

Such commitment devices could help people in some seemingly unlikely situations. As discussed above, Kaur et al. (2011) find that data-entry workers in India enthusiastically adopt a payment scheme that essentially penalizes them for not hitting targets. This idea could be applied to other situations where workers don’t work as hard as they themselves would like to, such as in factories or other informal work environments where workers are paid piece-rates but are unable to effectively monitor their own pace of work. Such commitment contracts could also help tackle the widespread problem of absenteeism among public-sector workers in developing countries such as India, where 25 per cent of government schoolteachers are absent from work on any given day, or Uganda, where 27 per cent are missing (Chaudhury et al. 2006).

Commitment devices are not the only way to help with self-control problems. We can also resolve the problem by allowing people to act on their good intentions *at the moment they have them*. One solution to the problem of low savings, for example, is to develop products that turn saving into an ‘impulse purchase’ by making it possible to buy savings at the store just as one might buy other products. One version of this idea, currently being tested in India by ideas42, involves selling stored-value cards such as the ones people routinely use to top up their mobile phones, except that the money spent on the ‘savings card’ adds to a person’s bank balance rather than his mobile air-time.

Principle 2: Reduce the need for self-control

It is true that we all have self-control problems. But it is also true that these problems are sometimes unintentionally created or made worse by the way policies are designed. A second way to tackle problems caused by a lack of self-control is therefore to find ways to reduce the need for people to exert self-control.

For example, poor families in the United States receive food stamps—transfers that enable them to buy food—at the beginning of the month. This system was designed for administrative convenience. But it worsens the effect of self-control problems: people over-spend in the first part of the month and are left with too little money for food by the end (see Gennetian et al. 2011). Switching from a single monthly installment to two fortnightly or four weekly ones in the case of food stamps would dramatically reduce the need for recipients to exert self-control in the first place.

Farmers face an even more challenging self-control problem. Harvest incomes come in once a season, sometimes once a year. Mani et al. (2013) show that sugarcane farmers for example pawn jewelry and consume a lot less in the months before harvest than in the months after harvest. This need not be the case. *Some* of the harvest income could easily be paid into an account that disburses a steady stream of monthly or fortnightly income instead of the one-off payout that is now common. The self-control problem is not inherent to the situation: changing the payout structure can rectify it.

Disaster-relief or other forms of compensation (such as compensation for being displaced by infrastructure or other projects) are also usually paid out in a single lump-sum. This imposes enormous self-control burdens on people who receive them. Switching to paying such benefits as a stream of payments over time rather than a single payment would mitigate these problems.

Developing countries, many of which are introducing or expanding cash transfers to the poor, can use these ideas to create more effective cash transfer programs than those in use in developed countries. For instance, income support programs for agricultural workers ought to be heavily loaded towards paying out in the agriculturally lean seasons. Using mobile-based or electronic methods of moving money allows the incorporation of more frequent transfers from the outset. This will help them avoid many of the problems which programs in the United States (such as the food stamps program discussed above) or Europe are only now beginning to try and rectify.

In many countries school fees and associated expenses are due at the beginning of a school year in one lump sum. This imposes self-control burdens, because many poorer parents cannot pay the entire year's fee out of their current income and so need to save up for it over time. Changing the timing of such fee demands to line up with the timing of income flows can increase people's ability to pay them. This problem is particularly marked for those with seasonal incomes (such as agricultural workers) whose income flows may not line up with the timing of such large payments. Switching from one large lump sum to allowing installment payments would also reduce the self-control burden.

Principle 3: Remove snags to choosing

We tend to assume that people make active decisions: faced with a set of options, they always actively choose the one they like most. However, behavioral economists have found that people frequently passively accept whatever happens if they do nothing. This means that the 'default option' is disproportionately important. Similarly, seemingly trivial steps, decisions, and choices (a form to fill or the need to submit one that has already been filled, for example) can drastically

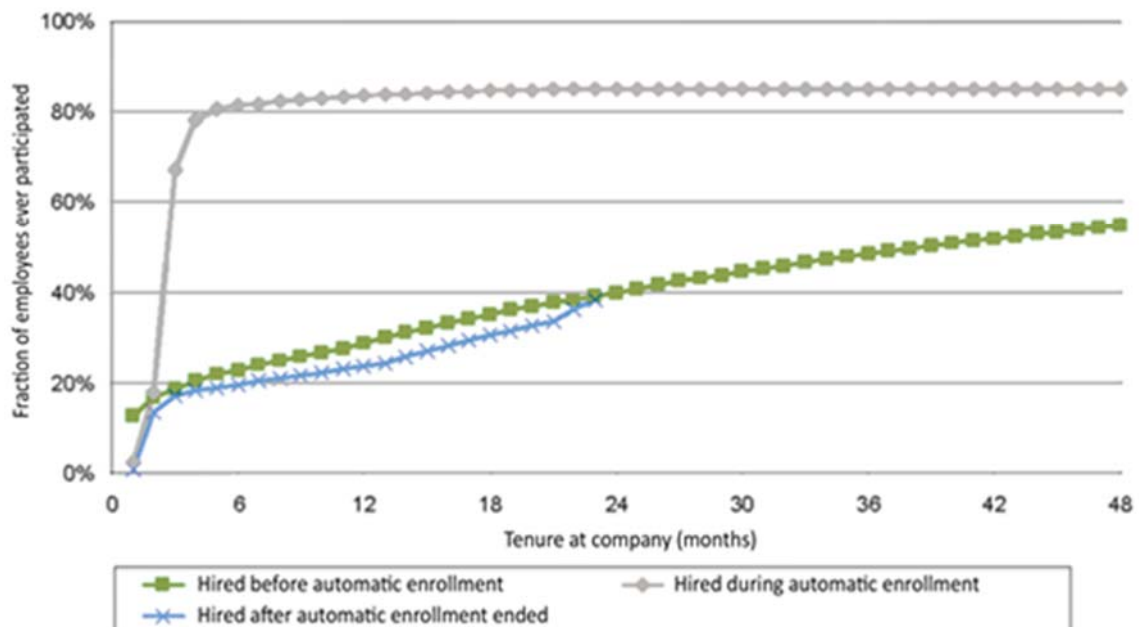
reduce the number of people who participate in a program. All this means that program uptake and use increases dramatically when the default is changed, or when a program is re-designed to reduce the number of things people have to do to take advantage of it.

In Morocco, Devoto et al. (2011) find that nearly 70 per cent of households who were helped with the administrative steps needed to get a piped water connection signed up for piped water, compared with just 10 per cent of those who were not. Reducing the hassle of participating in programs can thus have dramatic effects on how many people they reach.

Making things automatic also helps. Automatic transfers into savings accounts can increase saving rates by removing the small steps that stand between an intention to save and actual savings. Such automated transfers can be used to help people make the most of many kinds of income or transfer flows. Depositing a fraction of a benefit or crop payment into a savings account is far more likely to allow a farmer to avoid a cash shortage before the next payment comes in than allowing all of it to be close at hand. Workers in small or mid-sized firms in developing countries would benefit from having some of their income automatically put into savings accounts. Automation uses choice inertia—so often a source of forgone advantages—to people’s benefit.

The idea of manipulating defaults has revolutionized retirement savings programs in the United States. A decade ago, when employees in most American companies had to fill out a form in order to participate in their company’s 401(k) program and avail of the employer’s matching contributions, participation in 401(k)s was low. However, flipping the default option around from exclusion (i.e. having to actively ‘opt in’) to inclusion (i.e. having to ‘opt out’ if you don’t want to participate) raised participation in the retirement plan of the first test company from 37.4 per cent to over 85 per cent (Madrian and Shea 2001). This has been hugely influential: by 2006, 41.3 per cent of US employers with 5,000 or more employees used such a design, up from 17.3 per cent in 1999 (Profit Sharing/401(k) Council of America 2006).

Figure 6: Defaults and 401(k) participation



Source: Based on Madrian and Shea (2001).

A program for fortifying flour with iron, described in Banerjee and Duflo (2011) shows the power of this insight in a development context. The program was designed so that a household had to tell the miller whether it wanted to have its flour enriched only once; the miller was supposed to act accordingly each time they came back to him. Unfortunately, the participating millers flipped this around: they required the household to say whether they wanted iron added to their flour each time they brought grain to be milled. This changed default setting was enough to cause participation in the program to plummet, causing it to fail to achieve its objectives.

Unfortunately, unfavorable defaults are built into a number of features of poor people's lives. For example, we can rely on the water that is piped into our homes being safe because it has already been chlorinated. The poor on the other hand have to ensure that they use chlorine tablets if they want similarly clean water. Of course, they often forget to do this—as we would if we had to every time we filled drinking water—thus frustrating efforts to reduce deaths from diarrhea and other water-borne diseases.

Behavioral economists therefore argue that getting more people to use chlorine requires making its use as close to automatic as possible. A chlorine dispenser installed at the village well, dispensing exactly the right amount of chlorine each time at the press of a button, removes most of the steps that people normally have to take to chlorinate their water. In Kenya, Kremer et al. (2009) found that these dispensers were the most cost-effective way to reduce diarrhea that has so far been devised or tested.

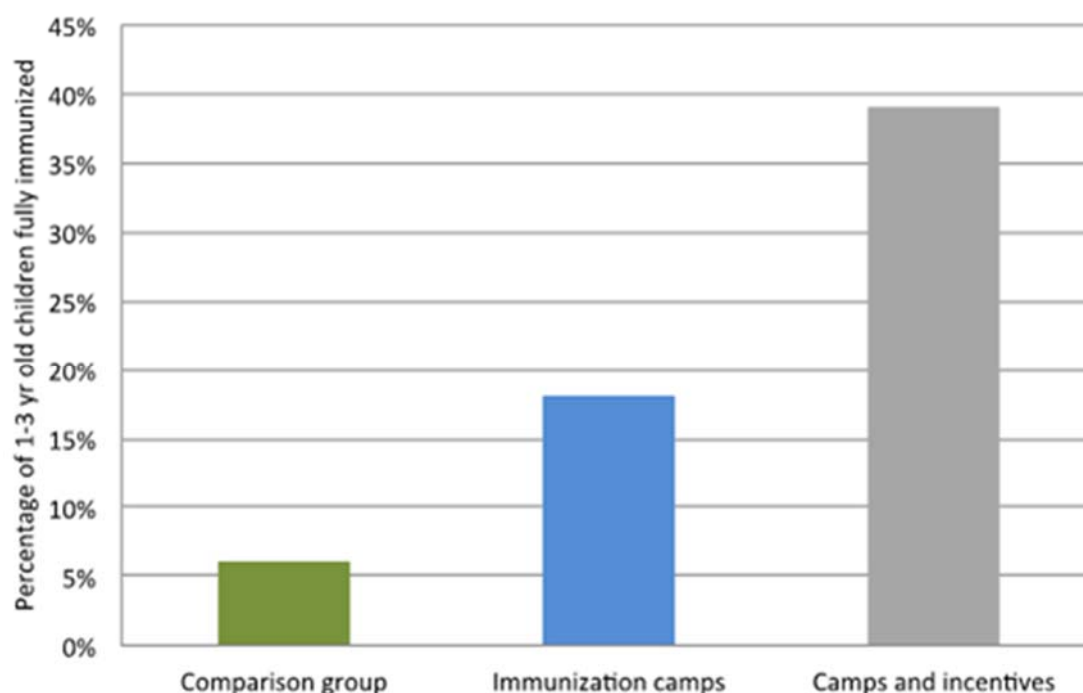
Similarly, many nutrition programs try either to get the poor to eat the kinds of food that naturally provide a balanced mix of micro- and macro-nutrients, or to get them to adopt special nutritional supplements. Most such programs have very little success. Behavioral science suggests that it might be most effective to make balanced nutrition close to automatic by fortifying food that people already eat with extra micronutrients, much as the routine iodization of salt has vastly reduced problems of iodine deficiency.

Principle 4: Use micro-incentives

We normally think that a *small* monetary or material incentive has no chance of inducing a change with *large* consequences. However, behavioral economists have found that such 'micro-incentives' affect how people behave in ways that have big consequences for their own well-being. The size of an incentive only needs to be as large as the barrier that caused the problem: if this is small, as it so often is, a small incentive is often enough.

For example, taking a child to a free immunization camp is a tiny inconvenience to endure for the protection provided by a full dose of vaccines. Yet by offering each parent who brought a child a half-kilo bag of lentils—equivalent to about half a day's wages for an agricultural laborer—succeeded in nearly doubling the fraction of children who were fully immunized in a part of rural India, from 18 to 29 per cent (Duflo et al. 2010). In addition, because nurses' wages, equipment costs and other program costs were fixed, adding the financial incentive actually resulted in halving the cost per child immunized.

Figure 7: Micro-incentives increase immunization




Source: Based on Banerjee et al. (2010).

In the United States, Volpp et al. (2008) offered patients on warfarin, an anti-stroke medication, a lottery ticket as a reward for taking their pills. Prior to the incentive, 20 per cent of the patients were not taking their medication correctly. The opportunity to win a small sum of money (the highest prize was \$100, which an individual had a 0.1 per cent chance of winning), however, succeeded in virtually eliminating incorrect drug adherence.

In Malawi, only 34 per cent of those getting tested for HIV at government-run testing centers were returning to pick up their results. Offering a tiny incentive of around \$0.15 (i.e. about 10 per cent of the daily wage) more than doubled the fraction of people who picked up their test results (Thornton 2008). Giving people a larger incentive—of up to \$3—did have a larger effect, raising the rate further to over 90 per cent. But the bulk of the jump—from 34 to over 70 per cent - was achieved simply by moving from no incentive to a tiny incentive.

Figure 8: Micro-incentives induced more people to pick up HIV test results



Source: Based on Thornton (2008).

Wider use of such micro-incentives could dramatically improve policies to control diseases like tuberculosis and HIV/AIDS (where medication is readily available but adherence is often a problem) as well as immunization in the developing world. The spread of mobile-phone technology opens up the possibility for innovative new ways of delivering such micro-payments. For instance, governments, mobile phone companies and pharmaceutical companies could collaborate on finding ways to transfer small amounts of airtime (or mobile cash) to people if they take their medicines.

Small incentives, frequently paid, could be used to reward a variety of ‘good behaviors’ in education such as school attendance. Recent evidence from the United States suggests that providing financial rewards for specific actions such as reading books or finishing homework raises test scores among students from low-income families (see Fryer 2010). Micro-incentives could also increase take-up of government programs that provide supplemental nutrition and other kinds of care to poor mothers or pregnant women, increasing the effectiveness of efforts to tackle problems like low birth weight and malnutrition.

Finally, more and more developing countries are moving to a point where the key constraint on achieving health and education goals is no longer building schools or clinics but rather ensuring that education and health workers show up to work. The World Bank estimates that 25 per cent of Indian schoolteachers, and 19 per cent of teachers in Bangladesh are absent from school on a given day (Chaudhury et al. 2006). Giving teachers or nurses a small incentive payment for each day over a certain minimum that they spend working could help tackle absenteeism, and maximize the impact of investments in health and education infrastructure.

Principle 5: Reduce inattention: reminders and implementation intentions

Behavioral economists have found that reminders—in person, using a phone call, or via text message, for example—can have dramatic positive effects on behaviors such as a failure to get tested for diseases, not taking medicines regularly, or even the tendency to incur penalties on high-interest borrowings. In all of these cases, following through on an intention requires a person to remember to take several steps, and it is easy to forget or neglect to do one of them. But missing

a single step often derails the whole process. A timely reminder goes a long way towards mitigating these problems.

For example, clients of three banks in Bolivia, Peru and the Philippines who were reminded (via letter or SMS) to make deposits into their accounts saved significantly more and were also more likely to reach a pre-set savings goal (Karlán et al. 2010). Similarly, Stango and Zinman (2011) find that having their attention drawn to their bank's policies about fees for overdrawn accounts reduces overdraft fees paid by individuals for up to two years after the reminders were sent.

Reminders have been used to increase workplace productivity. Cadena et al. (2011) found that reminding loan officers in a Colombian bank about their goals for credit disbursal, collections, etc. reduced their tendency to postpone contacting potential new clients or making efforts to collect on outstanding credit till just before their monthly bonuses were due. As a result, loan officers earned more and the quality of the bank's loan portfolio improved.

Encouraging people to make specific plans about when and how they will do something acts like a reminder, drawing attention to actions that might have been neglected. Such implementation intentions have been used to successfully encourage a number of health-related behaviors in the United States, including a 10 per cent increase in the number of people making and sticking to appointments to be screened for colon cancer (Milkman et al. 2012). Similarly, Milkman et al. (2011) show that nudging people to form plans about getting an influenza shot increased the fraction of people who got the flu vaccine.

The scope for applying the broad idea of using reminders and implementation intentions in developing countries is enormous. The spread of mobile phones makes it feasible to use text messages or calls to carry out the monitoring needed to ensure adherence to drugs for communicable diseases, reminding people to take their pills as near as possible to the actual times they need to take them. Timely reminders can be used to tackle other (non-medical) situations where people forgo significant benefits because they do not do something (such as weed their fields) at the right time.

Principle 6: Maximize the impact of messaging: framing effects, social comparisons, norms

Governments, agencies and non-governmental organizations communicate with the target audiences using information campaigns, billboards, letters, television or radio advertisements, and now personalized messaging through phones. Behavioral economics provides a number of principles about the content and framing of such messaging that can make it more effective at achieving its desired ends.

For example, messaging that links money with specific goals is an extremely effective way to increase savings rates. In the study of the effect of reminders on savings rates in Peru, Bolivia, and the Philippines discussed above (Karlán et al. 2010), reminders that emphasized a specific goal—a house, an appliance purchase, etc.—*were twice as effective* as ones that did not. This is because people treat money differently depending on what they think its purpose is. They are much more likely to hold off from spending a dollar associated with a longer-term goal than a dollar which, in their minds, is intended for general expenses.

People are much more responsive to being informed of what they lose by not doing something than they are to being told how much it benefits them. This insight can help refine the design of programs that seek to encourage people to take steps to ensure their own or their children's health: emphasizing the possible ill-effects from not vaccinating a child, for example, may be much more effective than emphasizing how healthy the child will be if she is vaccinated.

Thirdly, comparing a person to his peers, neighbors, friends, etc. is an extremely effective way to change behavior. For example, American households who got mailers that compared their own electricity consumption to that of homes in their own neighborhood reduced their power consumption by as much as they would have if the cost of power had risen by 11-20 per cent, with effects thrice as large for those who were initially using the most electricity (Allcott and Mullainathan 2010). Similarly large effects have been found for water consumption (Ferraro and Price 2011), where comparison to neighbors curbed water consumption far more than either simple information provision or messages exhorting people to be thrifty users of water. Drawing attention to what progressive neighbors are doing could thus spur the adoption of many beneficial technologies in developing countries.

Most individuals make efforts to conform to what they perceive the social norm to be. Sometimes, though, their perceptions about the norm may be inaccurate: less common behaviors may be more visible, leading people to think they are in fact 'what everyone does'. Messaging about social norms can change people's perceptions about what is normal and thus change behavior. This idea could help tackle a number of important problems in developing countries. For example, while 25 per cent of Indian teachers and 19 per cent of their Bangladeshi counterparts are missing from school each day, the fact remains that between 75 and 81 per cent do show up (Chaudhury et al. 2006). Being present is thus the norm, and emphasizing this may make those who routinely violate this norm less likely to do so. More girls now attend school than those who do not, even in countries where gender gaps in education persist. Drawing attention to this could help reduce these gaps further.

Finally, making a particular feature of a person, his environment, or a product more salient often has large effects on people's choices. For example, reminding a person (even inadvertently) of an aspect of their identity induces them to act in ways that fit in with the stereotypes associated with that aspect (Steele and Aronson 1995). A recent demonstration of this comes from India, where Hoff and Pandey (2004) gave village schoolchildren a set of simple puzzles to solve. They found that carrying on the experiment in a way that made it clear that participants' caste was known reduced the performance of lower-caste students dramatically, but did not affect the performance of others in mixed-caste groups. This suggests a need for carefully examining communications and publicity to ensure that they are not inadvertently strengthening damaging stereotypes or modifying them to evoke positive associations with aspects of people's identity.

Principle 7: Frame messages to match mental models

Existing mental models sometimes stand in the way of people adopting beneficial technologies or undertaking profitable investments in physical or human capital (see Section 3.4). For example, farmers who believe that fertilizer has no effect on productivity unless used in large quantities might forgo the proven effects of using even a small amount. The poor are also often unduly pessimistic about their own ability to affect outcomes, leading them to pass up on many productive investment opportunities. Frankenberger et al. (2007) found that a third of poor Ethiopian families believed that destiny was the single most important determinant of success. These families were less likely to make longer-term investments, making them less likely to ever escape poverty.

Carefully designed messaging can help in such situations. However, campaigns that simply tell people that their beliefs are inaccurate (for example, by reiterating that fertilizer increases productivity) are unlikely to make much headway. In part, this is because people usually disregard information that does not conform to their own mental model of a situation as irrelevant to their own circumstances ('It's not for me'). However, information or evidence that directly targets the beliefs at the core of the flawed mental model has a better chance of success.

The large gender gap in educational attainment between girls and boys in India is in part due to the belief that there are few economic benefits from educating girls, whose primary social function is believed to be limited to housework. Such beliefs reinforce themselves, since they lead parents to pull girls out of school, which in turn means that their economic opportunities in fact remain limited. But this problem is not unsolvable. Jensen (2012) finds that providing villagers with precise information about the availability of jobs for girls with high-school degrees and how to get such jobs causes teenage girls to stay in school longer, makes them more likely to look for paid work, and leads to them marrying later. Perhaps even more remarkable for an area with some of the worst gender gaps in education and health, primary-school-age girls in the villages which received the recruitment drives were 5 percentage points more likely to be in school and weighed more than in control villages. Parents had responded to the discovery that girls had economic prospects by investing more in their nutrition and education.

The case of adolescent sex education in Kenya provides another instructive lesson on the importance of tailoring messages to existing mental models. An education campaign in Kenya, where many teenaged girls were getting pregnant by older men, sought to reduce such pregnancies by urging girls to shun premarital sex. However, this reinforced the idea of marriage as a desirable goal, and girls viewed getting pregnant as the most efficient way to find a husband. The program therefore actually led teenaged girls to actively seek out older partners for unprotected sex. On the other hand, a campaign that simply provided girls with the information that older men were more likely to be HIV-positive reduced the number of girls who got pregnant by older men by two-thirds. It succeeded because it addressed the fundamental cause of such pregnancies, which was the perceived desirability of older men as sexual partners (Dupas 2011).

5.3 Testing and re-design

Which of these design principles is most useful in a particular situation will naturally depend on many factors. Often, there will be several possible paths forward. Identifying the interventions that seem most feasible and useful, rolling them out in a controlled way in a small pilot program, and tracking outcomes will inform an iterative process of re-design. Prototyping and a willingness to experiment and tweak are crucial. Over time, this will lead to an intervention or a small set of interventions that are both psychologically sound and administratively and logistically feasible. These can then be rigorously evaluated using the tools development economists have developed over the past decade, which compare the results of treatment groups to those of randomly chosen comparable control groups.

6 Conclusion

This paper is both a review of what is known about human behavior and how these insights have been applied to development policy, and an attempt to show what is possible if these insights and design principles are applied more broadly. We have tried to show how behavioral insights allow us to understand why many of the kinds of problems policymakers in developing countries face exist and persist, as well as allow us to design innovative, effective solutions to those problems.

We conclude by looking ahead to how these lessons can be applied in practice, and what this means for the way donors, researchers, and governments work. Adequately unlocking the potential of behavioral solutions will require us to take a systematic approach to identifying key problems, evaluating the potential impact of behavioral economics approaches, and translating these insights into improvements in programs. This can only be achieved by making some deep-seated changes in the way we go about applying behavioral insights in development.

First, efforts to apply behavioral economics insights have to be built around the objective of achieving impact at scale. This means moving away from a focus on relatively narrowly conceived research projects and ‘boutique’ pilots that aim to pin down a specific behavioral insight or insights and towards a focus on existing programs or projects that seek to address big development problems, but whose effectiveness is constrained by behaviors.

Second, innovation has to be embedded into the process of intervention design from the very beginning and must run all the way through it. The goal must not be to test one or two interventions but to design (possibly several) interventions based on careful problem analysis and the identification of behavioral bottlenecks. This process of design should be an iterative process that incorporates feedback from small tests carried out as part of the design itself.

Adopting this systematic approach towards diagnosis is important not just because it leads to better solutions to the problem in question but because it provides us with diagnoses and diagnostic techniques that can carry over to other contexts. Thus, if we find good diagnostics that indicate that self-control plays a large role in understanding a particular behavior in one country, we would have good reason to explore the use of these diagnostics in a different setting. Insights and diagnoses are likely to have external validity even if particular designs do not. In that sense, we should think of the kind of policy experimentation being described here partly also as mechanism experiments (see Ludwig et al. 2011).

For donors, this means selecting projects where successes can be scaled. This affects the kinds of problems chosen for experimentation. The most useful problems to work on would affect people in countries or regions beyond the one initially chosen, for instance because they are pervasive across a number of developing countries. There must also be reason to believe, *ex ante*, that behavioral barriers are a critical reason for program goals not being met. It also affects the choice of partners. Potential partners should reach large numbers of people so that any successes can be scaled up in the context of an existing program. Working with governments or large aid agencies may be more impactful than working directly with individual researchers. Involving private-sector players with established distribution and outreach networks can also be an effective way to scale, assuming that these players benefit from the behavioral problem being solved (whether directly in terms of profits or indirectly, by creating a base of consumers for other products).

For researchers, this focus on impact at scale means privileging projects that build on existing government or large-scale non-profit programs rather than collaborations with small, boutique NGOs or service providers. It also means being willing to evaluate an intervention that may not necessarily isolate the causal effect of a single psychology or pathway, but of a suite of linked design innovations. And it means paying close attention to the administrative burden or logistical requirements of any proposed solution, because these affect whether a solution can be scaled up.

Finally, this means that governments need to be open to involving behavioral experts when programs are first designed as well as to experimenting on existing programs. As important is openness to exploring new (and sometimes surprising) pathways to impact that emerge in the course of the detailed problem and behavioral analysis.

Embedding innovation into the design process itself leads to designs that have a greater chance of success than if we proceeded to testing the first feasible and reasonable set of ideas about how to solve a problem. Over time, a rigorous application of the approach to program design outlined in this paper should lead to more effective, cheaper and more easily replicable innovations. As we have seen, many policy problems can be traced in the ultimate analysis to gaps between intentions and actions. A systematic application of behavioral design should help close another, equally important gap: that between what policy seeks to achieve and what it accomplishes.

References

- Allcott, H., and S. Mullainathan (2010). 'Behavior and Energy Policy'. *Science*, 327(5970): 1204-5.
- Ashraf, N., D. Karlan, and W. Yin (2006). 'Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines'. *Quarterly Journal of Economics*, 121(2): 635-72.
- Banerjee, A., and E. Duflo (2011). *Poor Economics: A Radical Rethinking of the Way to Fight Global Poverty*. Public Affairs.
- Banerjee, A., E. Duflo, R. Glennerster, and D. Kothari (2010). 'Improving Immunisation Coverage in rural India: Clustered Randomised Controlled Evaluation of Immunisation Campaigns with and without Incentives'. *BMJ* 340:c2220.
- Banerjee, A., and S. Mullainathan (2008). 'Limited Attention and Income Distribution'. *American Economic Review*, 98(2): 489-93.
- Banik, P., A. Midya, B. K. Sarkar, and S. S. Ghose (2006). 'Wheat and Chickpea Intercropping Systems in an Additive Series Experiment: Advantages and Weed Smothering'. *European Journal of Agronomy*, 24: 325-32.
- Bernheim, B. D., D. Ray, and S. Yeltekin (2013). 'Poverty and Self-control'. Mimeo. New York: University Department of Economics.
- Bernheim, B. D., D. Ray, and S. Yeltekin (1999). 'Self-control, Savings, and the Low Asset Trap'. Mimeo. New York: New York University, Department of Economics.
- Bertrand, M., D. Karlan, S. Mullainathan, E. Shafir, and J. Zinman (2010). 'What's Advertising Content Worth? Evidence from a Consumer Credit Marketing Field Experiment'. *The Quarterly Journal of Economics*, 125(1): 263-305.
- Bertrand, M., S. Mullainathan, and E. Shafir (2004). 'A Behavioral Economics View of Poverty'. *The American Economic Review*, 94(2): 419-23.
- Brune, L., X. Gine, J. Goldberg, and D. Yang (2011). 'Commitments to Save: A Field Experiment in Rural Malawi'. Policy Research Working Paper 5748. Washington, DC: World Bank.
- Cadena, X., A. Schoar, A. Cristea, and H. Delgado-Medrano (2011). 'Fighting Procrastination in the Workplace: An Experiment'. Working Paper 16944. Cambridge, MA: NBER.
- Chaudhury, N., J. S. Hammer, M. Kremer, K. Muralidharan, and F. H. Rogers (2006). 'Missing in Action: Teacher and Health Worker Absence in Developing Countries'. *Journal of Economic Perspectives*, 20(1): 91-116.
- Ciriolo, E. (2011). 'Behavioural Economics in the European Union: Past, Present and Future'. Oxford: OXERA.
- Devoto, F., E. Duflo, P. Dupas, W. Pariente, and V. Pons (2011). 'Happiness on Tap: Piped Water Adoption in Urban Morocco'. *American Economic Journal: Public Policy*, 4(4): 68-99.
- Drexler, A., G. Fischer, and A. Schoar (2011). 'Keeping It Simple: Financial Literacy and Rules of Thumb'. *American Economic Journal: Applied Economics*, 6(2): 1-31.
- Duflo, E., M. Kremer, and J. Robinson, J. (2011). 'Nudging Farmers to Use Fertilizer: Evidence from Kenya'. *American Economic Review*, 101(6): 2350-90.
- Duflo, E., M. Kremer, and J. Robinson (2010). 'Why Don't Farmers Use Fertilizer? Evidence from Field Experiments in Kenya'. Working Paper. Published 2011 in *American Economic Review*, 101(6): 2350-90.

- Dupas, P. (2011). 'Do Teenagers Respond to HIV Risk Information? Evidence from a Field Experiment in Kenya'. *American Economic Journal: Applied Economics*, 3(1): 1-34.
- Ferraro, P., and M. Price (2011). 'Using Non-Pecuniary Strategies to Influence Behavior: Evidence from a Large Scale Field Experiment'. Working Paper 17189. Cambridge, MA: NBER.
- Frankenberger T. R., P. Sutter, A. Teshome, A. Aberra, M. Tefera, M. Tefera, A. S. Taffesse, T. Bernard, T. Spangler, and Y. Ejigsemahu (2007). 'Ethiopia: The Path to Self-Resiliency'. Report prepared for Partners in Rural Development on behalf Canadian Network of NGOs in Ethiopia.
- Fryer, R. G., (2010). 'Financial Incentives and Student Achievement: Evidence from Randomized Trials'. Published 2011 in *The Quarterly Journal of Economics*, 126(4): 1755-98.
- Gennetian, L., R. Seshadri, N. Hess, A. Winn, and R. George (2011). 'Running Out and Acting Out: Food Stamp Benefit Cycles and School Disciplinary Events among Chicago Public School Students'. Working Paper. Chicago, IL: University of Chicago.
- Hanna, R., S. Mullainathan, and J. Schwartzstein (2012). 'Learning Through Noticing: Theory and Experimental Evidence in Farming'. Working Paper 18401. Cambridge, MA: NBER.
- Hoff, K., and P. Pandey, (2004). 'Belief Systems and Durable Inequalities: An Experimental Investigation of Indian Caste'. Policy Research Working Paper. Washington, DC. World Bank.
- Jensen, R. (2012). 'Do Labor Market Opportunities Affect Young Women's Work and Family Decisions? Experimental Evidence from India'. *Quarterly Journal of Economics*, 127(2): 753-92.
- Kahneman, D. (2011). *Thinking, Fast and Slow*. Basingstoke: Macmillan.
- Karlan, D., M. McConnell, S. Mullainathan, and J. Zinman (2010). 'Getting to the Top of Mind: How Reminders Increase Savings'. Working Paper 16205. Cambridge, MA: NBER.
- Kaur, S., M. Kremer, and S. Mullainathan (2011). 'Self-Control at Work'. Working Paper. Cambridge, MA: MIT Sloan School of Management.
- Kremer, M., E. Miguel, S. Mullainathan, C. Null, and A. Zwane (2009). 'Making Water Safe: Price, Persuasion, Peers, Promoters, or Product Design?' Working Paper. Cambridge, MA: NBER.
- Ludwig, J., J. Kling, and S. Mullainathan, (2011). 'Mechanism Experiments and Policy Evaluations'. Working Paper 17062. Cambridge, MA: NBER.
- Madrian, B., and D. Shea (2001). 'The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior'. *The Quarterly Journal of Economics*, 116(4): 1149-87.
- Mani, A., S. Mullainathan, E. Shafir, and J. Zhao (2013). 'Poverty Impedes Cognitive Function'. *Science*, 341: 976-80.
- Milkman, K., J. Beshears, J. Choi, D. Laibson, and B. Madrian (2012). 'Following Through on Good Intentions: The Power of Planning Prompts'. Working Paper 17995. Cambridge, MA: NBER.
- Milkman, K., J. Beshears, J. Choi, D. Laibson, and B. Madrian (2011). 'Using Implementation Intentions Prompts to Enhance Influenza Vaccination Rates'. *Proceedings of the National Academy of Sciences*, 108(26): 10415-20.
- Morris, M., V. Kelly, R. Kopicki, and D. Byerlee (2007). 'Fertilizer Use in African Agriculture: Lessons Learned and Good Practice Guidelines'. Washington, DC: World Bank.

- Mullainathan, S., and R. H. Thaler (2000). 'Behavioral Economics'. Working Paper 7948. Cambridge, MA: NBER.
- Profit Sharing/401(k) Council of America (PSCA) (2006). '50th Annual Survey of Profit Sharing and 401(k) Plans', cited in M. Soto and B. A. Butrica (2009). 'Will Automatic Enrollment Reduce Employer Contributions to 401(K) plans?'. Working Paper 2009-33. Boston: Boston College Center for Retirement Research.
- Shah, A. K., S. Mullainathan, and E. Shafir (2012). 'Some Consequences of Having Too Little'. *Science*, 338(6107): 682-5.
- Stango, V., and J. Zinman (2011). 'Limited and Varying Consumer Attention: Evidence from Shocks to the Salience of Penalty Fees'. Working Paper 17028. Cambridge, MA: NBER.
- Steele, C., and J. Aronson. (1995). 'Stereotype Threat and the Intellectual Performance of African Americans'. *Journal of Personality and Social Psychology*, 69(5): 797-811.
- Thornton, R. (2008). 'The Demand for, and Impact of, Learning HIV Status'. *American Economic Review*, 98(5): 1829-63.
- UK Cabinet Office (2012). 'Applying Behavioural Insights to Reduce Fraud, Error and Debt'. Working Paper.
- Vohs, K., R. F. Baumeister, B. J. Schmeichel, J. M. Twenge, N. M. Nelson, and D. M. Tice (2008). 'Making Choices Impairs Subsequent Self-Control: A Limited-Resource Account of Decision Making, Self-Regulation, and Active Initiative'. *Journal of Personality and Social Psychology*, 94(5): 883-98.
- Volpp, K., G. Loewenstein, A. Troxel, J. Doshi, M. Maureen Price, M. Laskin, and S. Kimmel (2008). 'A Test of Financial Incentives to Improve Warfarin Adherence'. *BMC Health Services Research*, 8: 272.