Poverty and Self-Control

B. Douglas Bernheim, Stanford University and NBER
Debraj Ray, New York University
Sevin Yeltekin, Carnegie Mellon University

February 2, 2011

1   Introduction

The absence of self-control is often cited as an important contributory cause of persistent poverty, particularly (but not exclusively) in developing countries. Recent research indicates that the poor not only borrow at high rates,[1] but also forego profitable small investments.[2] To be sure, traditional theory — based on high rates of discount and minimum subsistence needs — can take us part of the way to an explanation. But it cannot provide a full explanation, for the simple reason that the poor exhibit a documented desire for commitment.[3] The fact that individuals are often willing to pay for commitment devices such as illiquid deposit accounts or ROSCA participation suggests that time inconsistency and imperfect self-control

---

[1] Informal interest rates in developing countries are notoriously high; see, for example Aleem (1990). But even formal interest rates are extremely high; for instance, the rates charged by microfinance organizations. Bangladesh recently capped formal microfinance interest rates at 27% per annum, a restriction frowned upon by the *Economist* (2010). Banerjee and Mullainathan (2010) cite other literature and argue that such loans are taken routinely and not on an emergency basis.

[2] Goldstein and Udrey (1999) and Udry and Anagol (2006) document high returns to agricultural investment in Ghana, even on small plots, while Duflo, Kremer, and Robinson (2010) identify high rates of return to small amounts of fertliizer use in Kenya, and de Mel, McKenzie, and Woodruff (2008) demonstrate high returns to microenterprise in Sri Lanka. Banerjee and Duflo (2005) cite other studies that also show high rates of return to small investments.

[3] See, for example, Shipton (1992) on the use of lockboxes in Gambia, Benartzi and Thaler (2004) on employee commitments to save out of future wage increases in the United States, Ashraf, Karlan, and Yin (2006) on the use of a commitment savings product in the Philippines, Aliber (2001), Gugerty (2001, 2007) and Anderson and Baland (2002) on ROSCA participation as a commitment device, and Duflo, Kremer, and Robinson (2010) on fertilizer use in Kenya (specifically, adoption just after harvest). See also the theoretical contributions of Ambec and Treich (2007) and Basu (2010).

are important explanations for low saving and high borrowing, complementary to those based on impatience, minimum subsistence or a failure of aspirations.

A growing literature already recognizes that the (in)ability to exercise self-control is central to the study of intertemporal behavior.[4] Our interest lies in how self-control and economic circumstances *interact*. If self-control (or the lack thereof) is a fixed trait, independent of personal economic circumstances, then the outlook for policy interventions that encourage the poor to invest in their futures – particularly one-time or short-term interventions – is not good. But another possibility merits consideration: poverty *per se* may damage self-control. If that hypothesis proves correct, then the chain of causality is circular, and poverty is itself responsible for the low self-control that perpetuates poverty.[5] In that case, policies that help the poor *begin* to accumulate assets may be highly effective, even if they are temporary.

The preceding discussion motivates the central question of this paper: is there some *a priori* reason to expect poverty to perpetuate itself by undermining an individual's ability to exercise self-control? Our objective requires us to define self-control formally and precisely. The term itself implies an internal mechanism, so we seek a definition that does not reference any externally-enforced commitment devices. Following Strotz (1956), Phelps and Pollak (1968) and others, we adopt the view that self-control problems arise from time-inconsistent preferences: the absence of self-control is on display when an individual is unable to follow through on a desired plan of action. What then constitutes the *exercise* of self-control? We take guidance from the seminal work of the psychologist George Ainslee (1975, 1992), who argued that people maintain personal discipline by adopting private rules (e.g., "never eat dessert"), and then construing local deviations from a rule as having global significance (e.g., "if I eat dessert today, then I will probably eat dessert in the future as well"). Formally,

---

[4]See, for instance, Akerlof (1991), Ainslie (1992), Thaler (1992), Laibson (1997), or O'Donoghue and Rabin (1999). There are social aspects to the problem as well. Excess spending may be generated by discordance within the household (e.g., husband and wife have different discount factors) or by demands from the wider community (e.g., sharing among kin or community).

[5]Arguments based on aspiration failures generate parallel traps: poverty can be responsible for frustrated aspirations, which stifle the incentive to invest. See, e.g., Appadurai (2004), Ray (2006), Genicot and Ray (2009) and the recent *United Nations Development Program* Report for Latin America based on this methodology. However, this complementary approach does not generate a demand for commitment devices.

with time-inconsistent preferences, it is natural to model behavior as a subgame-perfect Nash equilibrium of a dynamic game played by successive incarnations of the single decision-maker.[6] For such a game, any equilibrium path is naturally interpreted as a personal rule, in that it describes the way in which the individual is supposed to behave. Moreover, history-dependent equilibria can capture Ainslee's notion that local deviations from a personal rule can have global consequences.[7] For example, in an intrapersonal equilibrium, an individual might correctly anticipate that violating the dictum to "never eat dessert" will trigger an undesirable behavioral pattern. Under that interpretation, the scope for exercising self-control is sharply defined by the set of outcomes that can arise in subgame-perfect Nash equilibria.

We assume that time-inconsistency arises from *quasi-hyperbolic discounting* (also known as $\beta\delta$-discounting), a standard model of intertemporal preferences popularized by Laibson (1994, 1996, 1997) and O'Donoghue and Rabin (1999). To determine the full scope for deliberate self-control, we study the features of the set of *all* subgame-perfect Nash equilibria. To avoid excluding any viable personal rules, we impose no restrictions whatsoever on strategies (we do not require stationarity, for instance, or that the decision-maker punish deviations by reverting to the Markov-perfect equilibrium). This approach contrasts with the vast majority of the existing literature, which focuses almost exclusively on Markov-perfect equilibria (which allow only for payoff-relevant state-dependence), thereby ruling out virtually all interesting personal rules.[8] By studying the entire class of subgame-perfect Nash equilibria, we can determine when an individual can exercise sufficient self-control to accumulate greater wealth, and when she cannot. In particular, we can ask whether self-control is more difficult when initial assets are low, compared to when they are high.

The model we use is standard. There is a single asset which can be accumulated or

---

[6] This approach is originally due to Strotz (1956).

[7] This interpretation of self-control has been offered previously by Laibson (1997), Bernheim, Ray, and Yeltekin (1999), and Benhabib and Bisin (2001). See Bénabou and Tirole (2004) for a somewhat different interpretation of Ainslee's theory.

[8] Exceptions include Laibson (1994), Bernheim, Ray, and Yeltekin (1999), and Benhabib and Bisin (2001).

run down at some fixed rate of return. By using suitably defined present values, all flow incomes are nested into the asset itself. The core restriction is a strictly positive lower bound on assets, to be interpreted as a credit constraint. In other words, the individual cannot instantly consume *all* future income. The lower bound may therefore be interpreted as referring to that fraction of present-value income which she cannot currently consume.

Apart from this lower bound, the model is constructed to be scale-neutral. We assume that individual payoff functions are homothetic, so we deliberately eliminate any preconceived relationship between assets and savings that is dependent on preferences alone. (We return to this point when connecting our model to related literature.) Discounting is quasi-hyperbolic.

It is notoriously difficult to characterize the set of subgame-perfect Nash equilibria (or equilibrium values) for all but the simplest dynamic games, and the problem of self-control we study here is, alas, no exception. We therefore initially examined our central question by solving the model numerically using standard tools. Figure 1 illustrates the results of one such exercise (which we explain at greater length later in the paper). The horizontal axis measures assets in the current period, and the vertical axis measures assets in the next period. Thus, points above, on, and below the 45 degree line indicate asset accumulation, maintenance, and decumulation respectively. In this exercise, there is an asset threshold below which *all* equilibria lead to decumulation. Thus, with low assets, it is impossible to accumulate assets by exercising self-control through any viable personal rule; on the contrary, assets necessarily decline until the individual's liquidity constraint binds. In short, we have a poverty trap.

However, above that threshold, there are indeed viable personal rules that allow the individual to accumulate greater assets. Moreover, as we will see later, the most attractive equilibria starting from above the critical threshold lead to unbounded accumulation.
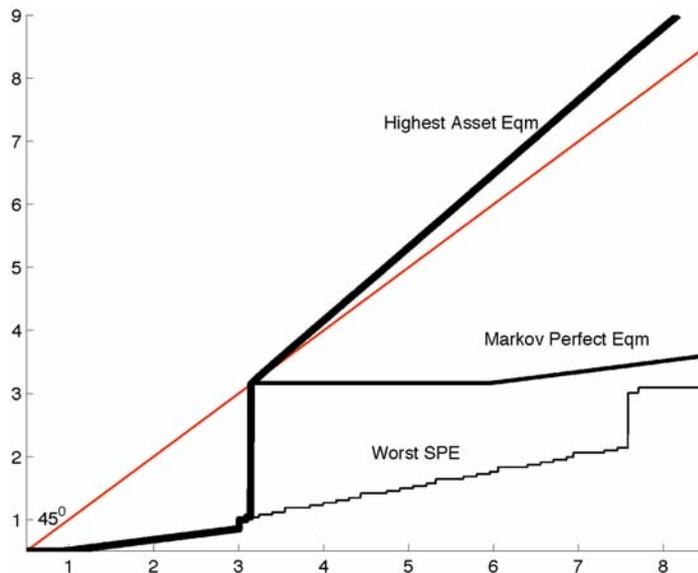
Possibilities for saving in subgame-perfect and Markov-perfect equilibria

The example suggests both our central conjecture and a (deceptively) simple intuitive argument in support of it. If imperfect capital markets impose limits on the extent to which an individual can borrow against future income, then potential intrapersonal "punishments" (that is, the consequences of deviating from a personal rule) cannot be all that bad when assets are low. If these mild repercussions are suitably anticipated, an individual will fail to exercise self-control. However, when an individual has substantial assets, she also has more to lose from undisciplined future behavior, and hence potential punishments are considerably more severe (in relative terms). So an individual would be better able to accumulate additional assets through the exercise of self-control when initial assets are higher. Obviously, if time inconsistency is sufficiently severe, decumulation will be unavoidable regardless of initial assets, and if it is sufficiently mild, accumulation will be possible regardless of initial assets (provided the individual is sufficiently patient). But for intermediate degrees of time inconsistency, we would expect decumulation to be unavoidable with low assets, and accumulation to be feasible with high assets.

It turns out, however, that the problem is considerably more complicated than this simple intuition suggests. (The overwhelmingly numerical nature of our earlier working paper,

5

Bernheim, Ray, and Yeltekin (1999), bears witness to this assertion.) The credit constraint at low asset levels infects individual behavior at *all* asset levels. In particular, they affect the structure of "worst personal punishments" in complex ways, as assets are scaled up. The example of Figure 1 illustrates this point quite dramatically: there are asset levels at which the *lowest* level of continuation assets jumps up discontinuously. As assets cross those thresholds, the worst punishment becomes *less* rather than *more* severe, contrary to the intuition given above. Thus, on further reflection, it is not at all clear that the patterns exhibited in Figure 1 will emerge more generally.

Our main theoretical result demonstrates, nevertheless, that the central qualitative properties of Figure 1 are quite general. For intermediate degrees of time inconsistency (such that accumulation is feasible from some but not all asset levels), there is a threshold asset level below which accumulation is impossible, and above which decumulation is avoidable. There is always an asset level below which liquid wealth is exhausted in finite time (and hence a poverty trap), as well as a level above which the most attractive equilibria give rise to unbounded accumulation.

One might object to our practice of examining the entire set of subgame-perfect equilibria on the grounds that many such equilibria may be unreasonably complex. On the contrary, we show that worst punishments have a surprisingly simple "stick-and-carrot" structure:[9] following any deviation from a personal rule, the individual consumes aggressively for one period, and then returns to an equilibrium path that maximizes her (equilibrium) payoff *exclusive* of the hyperbolic factor. Thus, all viable personal rules can be sustained without resorting to complex forms of history-dependence.

In addition, it is unclear that Markov equilibria in this model are any less "complex", despite their "simple" dependence on just the payoff-relevant state. As we shall show, Markov equilibria typically involve several jump discontinuities, and their payoff behavior as

---

[9]Though there is an obvious resemblance to the stick-and-carrot punishments in Abreu (1988), the formal structure of the models and the arguments differ considerably. Most obviously, Abreu considered simultaneous-move repeated games, rather than sequential-move dynamic games with state variables.

a function of initial assets, suitably normalized, is typically nonmonotonic. Also, identifying Markov equilibria is more computationally challenging than determining the key features of subgame-perfect equilibria. In any case, we demonstrate that, when consumers face liquidity constraints, Markov-perfect equilibria also give rise to asset traps.

Our analysis has a number of provocative implications for economic behavior and public policy. First (and most obviously), the relationship between assets and self-control argues for the use of "pump-priming" interventions that encourage the poor to start saving, and rely on self-control to sustain frugality at higher levels of assets. Second, policies that improve access to credit (thereby relaxing liquidity constraints) reduce the level of assets at which asset accumulation becomes feasible, thereby helping more individuals to become savers (although those who fail to make the transition fall further into debt). Intuitively, with greater access to credit, the consequences of a break in discipline become more severe, and hence that discipline is easier to sustain to begin with. Third, the opportunities to make commitments may be significantly less valuable to those with self-control problems than previous analyses have implied. For example, in certain circumstances, individuals with self-control problems will avoid opportunities to lock up funds (e.g., in retirement accounts or fixed deposit schemes), even when they wish to save. This occurs when desired saving exceeds the maximum amount that can be locked up, but not by too much. In such cases, locking up funds moderates the consequences of a lapse in discipline, thereby making self-control more difficult to sustain. Finally, we argue that the model can potentially provide an explanation for the observed "excess sensitivity" of consumption to income.

As noted above, we build on an earlier unpublished working paper (Bernheim, Ray, and Yeltekin (1999)), which made our main points through simulations, but did not contain theoretical results. Our work is most closely related to that of Banerjee and Mullainathan (2010), who also argue that self-control problems give rise to low asset traps. Though the object of their investigation is similar, their analysis has little in common with ours. They examine a novel model of time-inconsistent preferences, in which rates of discount differ from

one good to another. "Temptation goods" (those to which greater discount rates are applied) are inferior by assumption; this assumed non-hometheticity of preferences automatically builds in a tendency to dissave when resources are limited, and to save when resources are high.

It is certainly of interest to study poverty traps by hardwiring non-homothetic self-control problems into the structure of preferences. Whether a poor person spends proportionately more on temptation goods than a rich person (alcohol versus IPods, say) then becomes an empirical matter. But we avoid such hardwiring entirely by studying homothetic preferences in an established model of time-inconsistency. The phenomena we study are traceable to a single built-in asymmetry: an imperfect credit market. Every scale effect in our setting arises from the interplay between credit constraints and the incentive compatibility constraints for personal rules. The resulting structure, in our view, is compelling in that it requires no assumption concerning preferences that must obviously await further empirical confirmation. In summary, though both theories of poverty traps invoke self-control problems, they are essentially orthogonal (and hence potentially complementary): Banerjee and Mullainathan's analysis is driven by assumed scaling effects in rewards, while ours is driven by scaling effects in punishments arising from assumed credit market imperfections.[10]

The rest of the paper is organized as follows. Section 2 presents the model. Section 3 examines Markov perfect equilibria, while Section 4 studies all subgame-perfect equilibria. Section 5 presents conclusions and some directions for future research.

---

[10]While our model is also related to Laibson (1994) and Benhabib and Bisin (2001), our analysis is not. These two papers consider history-dependent strategies in a *fully* scalable model, in which both preferences are homothetic and there is no credit constraint. It follows, as we observe more formally below, that every equilibrium path can be replicated, by scaling, at all levels of initial assets, so that there is no relationship between poverty and self-control.