

Climbing Out of Poverty: Long Term Decisions under Income Stress

Abhijit V. Banerjee

Sendhil Mullainathan
MIT and Harvard University

August 2007

1 Introduction

There is now a broad consensus that self-control problems are very important in understanding diverse phenomena ranging from the economics of health clubs to the demand for pension plans. This paper attempts to investigate in what ways thinking about self-control issues can help us understand better the behavior of the poor and the special challenges they face.

A recent literature has emphasized that there are lots of gaps in our understanding of the poor. In particular we do not understand why the poor do not respond more to the availability of highly profitable and entirely divisible opportunities for investment. For example, Duflo, Kremer and Johnson (2007) point out the adoption of fertilizer by Kenyan farmers is puzzlingly slow, given how profitable it has been shown to be, and the fact that it can be used even on very small plots of land. Similar points have been made about the planting of pineapples in Ghana (Conley and Udry, 2005) and planting trees (IFC, XXXX). Banerjee and Duflo (2006) is another paper that points out several seemingly anomalous behavior patterns that we observe among the poor, including the fact that many of them seem to have a quite bit of slack in their budgets, but do not use it to pay down their often prohibitively expensive debt.

One explanation of much of this is that the poor are simply impatient. Indeed, there is long tradition in economics of assuming that the poor are poor precisely because they are impatient. The problem with this explanation is that the same poor people who fail to make these investments often participate in ROSCAs, which typically offer negative returns.

In this paper, we argue that a simple model of temptations has a lot to offer in terms of understanding why the poor behave the way they do. In particular we develop a model where the degree of "temptedness" is endogenous and might vary systematically between the rich and the poor. We argue that this formulation helps resolve many of the puzzles and moreover raises the possibility of a behavioral poverty trap, where people are present-biased because they are

poor, but that in turn keeps them poor. In other words the impatience that the poor often shown is as much a result of their poverty as it is a cause.

The key to these results is that the assumption that consumption has two components: x_t and z_t . The first, x_t , reflects consumption on which there is no temptation. Both long-run and short-run selves value this consumption. The second, z_t , reflects consumption where there is temptation—the utility the consumer gets from z is only valued by the self that is taking the decision—future selves do not value the fact that you have given into a temptation. We show that this is a simple generalization of the hyperbolic model. Yet this generalization makes it much more useful for the understanding the poor, because it allows us to consider how the “shape” of temptation changes as incomes rise, a point we return to below.

The core technical result from this model which makes it tractable is that we can derive a modified Euler equation to describe consumption. In the traditional Euler equation individuals would equate marginal utility today ($u(x)$) with the expected discount-rate and interest-rate adjusted marginal utility tomorrow ($\delta ru(x)$). In the temptation-modified Euler equation, the marginal utility tomorrow is further adjusted by a “temptation tax”. Individuals know that a dollar transferred to tomorrow would be partly spent on x goods and partly spent on z goods. Since from the vantage point of today, only the x goods “matter”, the z consumption can be viewed as a tax. So marginal utility today is equated to $\delta Ru(x)[1 - z'(c)]$ where $z'(c)$ is the marginal propensity to consume z .

The rest of the paper after developing the modified Euler equation constitutes of a number of examples that bring out the implications of this temptation tax. We first show that it is a very specific kind of tax, and in particular one that has particularly negative effects on savings. This idea is then generalized, in the sense that we show that uncertainty about temptations can also hurt savings.

Next we look at what happens when the temptation tax depends on your level of consumption. To see what this might buy us, consider a two period model where the tax, $z(c)$, is concave. Individuals who are poor would face a steep $z'(c)$ and therefore be less willing to save. They realize that any marginal dollar they save will likely be spent on temptation goods that they do not save. On the other hand, consider a wealthier individual. His higher overall level of consumption implies that he will be on a steeper part of the $z(c)$ and therefore is in better position to save. Notice that these results arise without either lumpy investments or credit constraints. They arise out of the inherent nonlinearity of temptation.

With multiple periods, this effect generalizes into a poverty trap. We show specific conditions on the shape of $z(\cdot)$ that produce a poverty trap: individuals below a wealth \tilde{w} will not save and will tend towards a low steady state level of wealth. Individuals above wealth \tilde{w} will save and tend towards a high steady state level of wealth. Conversely, we show that if the temptation function is convex, this type of poverty trap cannot occur. The intuition for these results is a dynamic generalization of the two period model. Individuals realize that if

their wealth gets high enough the temptation tax lowers. This gives incentives to low wealth individuals to save in the hopes of reaching this high wealth-low temptation equilibrium. If, however, they are sufficiently poor it will take them too long to get to this low tax wealth level and therefore they will give up.

Finally, we argue that this particular shape of the temptation function, namely that $z(c)$ is concave can help us resolve a number of puzzles about the behavior of the poor

2 Overview of Model

We model utility as having two parts:

$u(x)$: x goods represent goods where there are no momentary temptations

$v(z, \zeta)$: z goods represent “temptations”. We introduce ζ into the preference to allow for random shocks to preferences (you might see something and get tempted to buy it).

The way to see the distinction between x goods and z goods is to look at the forward-looking preferences:

$$u(x_0) + v(z_0) + E \sum_1^T \delta^t u(x_t)$$

where E is the expectations operator.

Note that while today’s self cares about both x and z consumption, he does not care about the future selves’ consumption of z . Define $c_t = x_t + z_t$

We assume that the household begins with wealth w_0 , has no labor income and faces an inter-temporal budget constraint:

$$w_{t+1} = f(w_t - c_t, \theta)$$

θ represents a shock to the production function. Let $G(\xi, \theta)$ represent the joint distribution of ξ and θ . We assume for the time being that the household cannot

borrow and that the sole savings technology is embodied in $f(\cdot)$.

The fact we are not taking into account all the future utility ought to be reminiscent of the hyperbolic discounting models and indeed we can generate a hyperbolic utility function from this formulation if we set

$$\begin{aligned} z_t &= x_t = c_t \\ u(x_t) &= \beta \tilde{u}(c_t) \\ v(z_t) &= (1 - \beta) \tilde{u}(c_t) \end{aligned}$$

In other words we are imposing the extra restriction that the person cannot separately choose x and z : Instead he chooses a single c_t . However then he only cares about a part of that c consumption in the future. These assumptions give us the hyperbolic maximand

$$\tilde{u}(c_0) + \sum \beta \delta^t \tilde{u}(c_t).$$

Perhaps a more intuitive way to look at the connection between the hyperbolic model and our formulation, is to set

$$\begin{aligned} u(x_t) &= \frac{x_t^{1-\alpha}}{1-\alpha} \text{ and} \\ v(z_t) &= A \frac{z_t^{1-\alpha}}{1-\alpha} \end{aligned}$$

In other words two essentially identical hyperbolic functions: Under these assumptions, within period choice between x and z will give us $z_t = qx_t$. Substituting this into our maximand gives us

$$\begin{aligned} & \frac{x_0^{1-\alpha}}{1-\alpha} + \frac{(q)^{1-\alpha} x_0^{1-\alpha}}{1-\alpha} + \sum \delta^t \frac{x_t^{1-\alpha}}{1-\alpha} \\ &= (1+q^{1-\alpha}) \left[\frac{x_0^{1-\alpha}}{1-\alpha} + \sum \beta \delta^t \frac{x_t^{1-\alpha}}{1-\alpha} \right] \end{aligned}$$

where $\beta = \frac{1}{1+q^{1-\alpha}}$, which is exactly in the hyperbolic form.

However for more general preferences the two models are quite distinct. The hyperbolic model can be written as

$$\frac{1}{\beta} u(x_0) + \sum \delta^t u(x_t)$$

while our model can be written as:

$$\begin{aligned} & u^*(x_0) + \sum \delta^t u(x_t) \\ & \text{where} \\ u^*(x_0) &= u(x_0) + v(z(x_0)) \end{aligned}$$

There is an additional wedge between u and u^* function is key to capturing a set of additional factors that we want to emphasize. In particular, because

$$u^*(x) = u(x)[1 + z(x)]$$

what will be key in our analysis of the model is the shape of the $z(x)$ function. However the $z(x)$ function is derived from the equation

$$u(x) = v(z(x)).$$

Therefore by defining $f(z)$ such that $f(z(x)) = x$ (i.e, the inverse function of $z(x)$), and

$$v(z) = \int_0^z u(f(z)) dz,$$

we can generate any $z(x)$ function we want. One important special case is $z = ax$. To get this all we need is to assume

$$v(z) = au\left(\frac{1}{a}z\right).$$

2.1 The Modified Euler Equation

These preferences set up a game between the present self and subsequent selves. We will be looking for a sub-game perfect equilibrium of this game in Markovian Strategies with household wealth and the contemporary realizations of the state variable. This section shows that for the infinite horizon case it is possible to characterize the equilibrium by a relationship similar to the traditional Euler Equation for intertemporal consumption optimization. Since wealth and the

Several points are worth noting about this Modified Euler Equation.

First, it is an inter-temporal equation. It can be combined with the intra-temporal equation:

$$u(x_t) = v(z(x_t))$$

to define the consumption path.

Second, as we will see in Section 3, this model can produce $x(w_t)$ that is non-differentiable. We show in the Appendix that a more general "equivalent" of the Modified Euler Equation holds even in these cases.

Finally, it bears close resemblance to the traditional Euler equation that would result if $v(z_t) = 0$:

$$u(x_t) = \delta f(w_t) u(x_{t+1})$$

The key difference is the added term $[1 - z_x(x_{t+1})]$ which serves as a "tax" on

future consumption. In other words individuals do not equalize marginal utilities over time. Instead tomorrow's marginal utility is taxed at a rate equal to the future self's propensity to consume frivolously. Put another way, any extra dollar of x consumption tomorrow is accompanied by $z(x)$ dollars of frivolous consumption which is unvalued by today's self. The key insight of the Modified Euler Equation is that the shape of the $z(x)$ determines the consumption path.

A simple example might clarify the role of $z(x)$: Compared to the case where $\frac{z(x)}{c}$ is a constant, the case where $\frac{z(x)}{c}$ is declining in x gives people an added reason to accumulate wealth—in order to reduce frivolous consumption. In the next section we take up some more structured examples which help us understand better the effect of temptations on wealth accumulation.

3 Some Examples

3.1 Example 1: The CRRA case

Suppose that utility is defined by two CRRA functions. So utility is $u(x) = \frac{x^{1-\alpha}}{1-\alpha}$ and $v(z) = \frac{z^{1-\alpha}}{1-\alpha}$ and $a = q^\alpha$ so that instantaneous utility at time t is $\frac{x_t^{1-\alpha}}{1-\alpha} + q^\alpha \frac{z_t^{1-\alpha}}{1-\alpha}$. The intra-temporal relationship in this case implies that

$$\begin{aligned} u(x_t) &= v(z_t) \\ x_t^{1-\alpha} &= q^\alpha z_t^{1-\alpha} \\ qx_t &= z_t \end{aligned}$$

so wasted consumption z is simply a constant fraction of valued consumption x .

The intertemporal utility function in this case is given by

$$\frac{x_0^{1-\alpha}}{1-\alpha} + q^\alpha \frac{x_0^{1-\alpha}}{1-\alpha} + \sum_1 \delta^t \frac{x_t^{1-\alpha}}{1-\alpha}$$

which can be rewritten as

$$(1+q)\frac{x_0^{1\Box\alpha}}{1-\alpha} + \sum_1 \delta^t \frac{x_t^{1\Box\alpha}}{1-\alpha}.$$

In other words, this case collapses exactly to the case of hyperbolic discounting and a higher q is exactly equivalent to a more set of hyperbolic preferences. More generally however these are quite different ways to parametrize preferences, though both involve time-inconsistent behaviors.

3.2 Example 2: Temptations and savings

In addition to the above preferences, assume that the production function is linear $w_{t+1} = R(w_t - c_t)$. Application of the Modified Euler equation to this case yields:

$$\begin{aligned} u(x_t) &= \delta f(w_t)u(x_{t+1})[1 - z(w_{t+1})] \\ x_t^{\Box\alpha} &= \delta R x_{t+1}^{\Box\alpha} [1 - qx(w_{t+1})] \end{aligned}$$

To solve this, assume that $c_t = x_t + z_t = \gamma w_t$, where the key constraint is that γ is a constant. Also denote by $\theta = \frac{q}{1+q}$, the ratio of z consumption in c_t . Substitution yields that γ is defined by

$$(1-\gamma)^\alpha = \delta R^{1\Box\alpha} (1-\gamma\theta), \quad (2)$$

so that a constant γ is consistent with a solution to the Euler equations.

To see what this implies, figure 1 plots the two sides of this equation as a function of γ . As long as we make the standard assumption that $\delta R^{1\Box\alpha} < 1$, the left-hand-side is greater than the right-hand-side at $\gamma = 0$, and as long as $\theta > 0$, the right-hand-side dominates at $\gamma = 1$. Given that the right-hand-side is linear while the left-hand-side is either everywhere concave or everywhere convex, the two curves must intersect once and only once. In other words, there is a single solution to the modified Euler equation.

The reason why this uniqueness property is not entirely obvious is that γ enters negatively on both sides of this equation. Essentially, as the Modified Euler Equation makes clear, the incentive to save today depends on the savings rate tomorrow, because a part of what gets consumed tomorrow will go to z and today's self only puts value on x consumption. Lower consumption tomorrow means less leakage and hence encourages saving today, which raises the possibility that we get multiple equilibria: people consume more today because they expect their future selves to also consume more. The standard assumption that $\delta R^{1\Box\alpha} < 1$ rules this out for the CRRA case. However there are preferences where this multiplicity might be of independent interest.

Since the curve representing $\delta R^{1\Box\alpha} (1-\gamma\theta)$ is figure 1 intersects the $(1-\gamma)^\alpha$ curve from below, an increase in θ (which moves it down) must increase the chosen γ . In other words, temptations reduce savings.

This is less obvious than it might seem. We have suggested that temptations act as a tax on future consumption from the point of view of today's self. However it is well-known that a tax on the return on capital does not always reduce savings because there is both a substitution effect and an income effect. Indeed it is evident from equation 2 and figure 1 that the effect of a reduction in R on γ can go either way, depending on whether $\alpha < 1$ or the other way around. Therefore temptations are actually a very particular type of tax, which happen to unambiguously (at least under the assumptions made) reduce savings.

3.3 Example 3: Temptation shocks and savings

The idea that people face a constant level of temptation is of course entirely implausible. Introspection suggests that temptation is more in the nature of a shock, and people go between being more or less susceptible to the attractions of whatever frivolous need that happens to attract them.

Before we look at the effects of temptation shocks, it is worth going back to the role of preference shocks in the standard model of savings as a useful benchmark. The first observation, which ought to be obvious is that if the demand for z is not a "temptation" in the sense that the decision maker actually cares about z consumption in the future, then the weight on u versus v (i.e. q) will not matter in the saving decision.

What about the effect of q varying when there are no temptations? The Euler equation for this case turns out to be

$$x_t^{\square\alpha} = \delta R E[x_{t+1}^{\square\alpha}]$$

which, using the intertemporal budget constraint and assuming, once again that there is a solution of the form $x_t = m(q)y_t$, can be written in the form

$$m(q)^{\square\alpha} = \delta R^{1\square\alpha} E[m(\tilde{q})^{\square\alpha} ((1 - m(q)(1 + q))^{\square\alpha})]$$

where q denotes the current (realized) q and \tilde{q} the anticipated random q . We can rewrite this in the form

$$\left[\frac{(1 - m(q)(1 + q))}{m(q)} \right]^{\alpha} = \delta R^{1\square\alpha} E[m(\tilde{q})^{\square\alpha}] = J_{NT}, \quad (3)$$

which implies

$$\begin{aligned} m(q) &= \frac{1}{1 + q + J_{NT}^{\frac{1}{\alpha}}} \\ \text{and} \\ \gamma(q) &= m(q)(1 + q) = \frac{1 + q}{1 + q + J_{NT}^{\frac{1}{\alpha}}}. \end{aligned}$$

Using this we can rewrite the above equation in the form

$$J_{NT} = \delta R^{1\square\alpha} E[1 + \tilde{q} + J_{NT}^{\frac{1}{\alpha}}]^{\alpha}.$$

Defining $\tilde{q} = \bar{q} + \hat{q}$ where \bar{q} is deterministic, and $\widetilde{J_{NT}} = \frac{J_{NT}^{\frac{1}{\alpha}}}{1+\bar{q}}$, we can rewrite the above equations in this more revealing form

$$\begin{aligned}\tilde{J}_{NT}^{\alpha} &= \delta R^{1-\alpha} E[1 + \frac{\hat{q}}{1+\bar{q}} + \widetilde{J_{NT}}]^{\alpha} \\ E\gamma(q) &= E[\frac{1 + \frac{\hat{q}}{1+\bar{q}}}{1 + \frac{\hat{q}}{1+\bar{q}} + \widetilde{J_{NT}}}] .\end{aligned}$$

What this equation makes clear is that as long as the distribution of $\frac{\hat{q}}{1+\bar{q}}$ is unchanged, γ cannot change. So bigger shocks do not affect savings as long as the distribution of $\frac{\hat{q}}{1+\bar{q}}$ is unchanged.

This neutrality result is the product of the fact that in the no-temptations case, the level of q does not matter, because it does not matter whether we are eating x or z . Variation in q does matter but that is being controlled by keeping the distribution of $\frac{\hat{q}}{1+\bar{q}}$ fixed.

Next, consider what happens in the same setting if the z s are temptations. The modified Euler Equation gives us that

$$m(q)^{\square\alpha} = \delta R^{1-\alpha} E[m(\tilde{q})^{\square\alpha} (1 - m(\tilde{q})\tilde{q}) ((1 - m(q)(1+q))^{\square\alpha})]$$

which can be rewritten as

$$[\frac{(1 - m(q)(1+q))}{m(q)}]^{\alpha} = \delta R^{1-\alpha} E[m(\tilde{q})^{\square\alpha} (1 - m(\tilde{q})\tilde{q})] = J_T.$$

Comparing this equation to equation 3, which is the corresponding equation for the previous case, we see that the left-hand-side is unchanged, but the right-hand-side, which represents the value of the future, has gone down for any fixed distribution of $m(\tilde{q})$, because $(1 - m(\tilde{q})\tilde{q})$ is less than 1. This tells us that for any realization of q , $m(q)$ today will be more when there are temptations than when there are none, assuming that the propensity to consume in the future has not gone up. But the propensity to consume in the future will also go up, precisely for this reason (i.e. the distribution of $m(\tilde{q})$ shifts up because it is the same as the distribution of $m(q)$). This compounds the tendency of $m(q)$ to go up, so the net result can be a larger increase in $m(q)$ compared to the no-temptations case.

By analogy to the previous case, it is easily shown that

$$\begin{aligned}m(q) &= \frac{1}{1+q+J_T^{\frac{1}{\alpha}}} \\ \text{and} \\ \gamma(q) &= m(q)(1+q) = \frac{1+q}{1+q+J_T^{\frac{1}{\alpha}}} .\end{aligned}$$

Both these equations are entirely unchanged. What changes is the equation determining J_T which, after some manipulation can be written in the form

$$J_T = \delta R^{1-\alpha} E[1 + \tilde{q} + J_T^{\frac{1}{\alpha}}]^{\alpha} [1 + J_T^{\frac{1}{\alpha}}] .$$

Using the same decomposition for \tilde{q} as before and defining \tilde{J}_T analogously to \tilde{J}_{NT} , we can write this in the form

$$\tilde{J}_T^\alpha = \delta R^{1-\alpha} E[1 + \frac{\hat{q}}{1 + \bar{q}} + \tilde{J}_T]^{\alpha-1} [\frac{1}{1 + \bar{q}} + \tilde{J}_T].$$

Clearly an increase in the size of the shocks which keeps the distribution of $\frac{\hat{q}}{1 + \bar{q}}$ unchanged will reduce \tilde{J}_T since $1 + \bar{q}$ goes up. From the expression for the consumption rate

$$E\gamma(q) = E[\frac{1 + \frac{\hat{q}}{1 + \bar{q}}}{1 + \frac{\hat{q}}{1 + \bar{q}} + \tilde{J}_T}]$$

it is evident that since the distribution of $\frac{\hat{q}}{1 + \bar{q}}$ is unchanged and \tilde{J}_T goes down, consumption will go up when the shocks get bigger. In other words, we have the following result:

Proposition 2 *An increase in the size of z -shocks that would have no effect on consumption in the no-temptation case, will reduce savings in the case where the z goods represent temptations.*

In other words, uncertainty about future *needs* tends to be more pro-savings than uncertainty about future (unjustifiable) *wants*.

3.4 Example 4: Non-homothetic preferences

All the examples so far assume that the preferences are homothetic. However the main advantage of this formulation is that it can easily accomodate different variants of non-homothetic preferences. In particular the pull of temptations may go up or down as people get richer. We now turn to the implications of such non-homotheticities.

We now introduce a generalization of the utility function we have been working with so far:

$$\begin{aligned} u(x) &= \frac{x^{1-\alpha}}{1-\alpha}, \\ a &= (q^b)^\alpha \end{aligned}$$

and

$$v(z) = \begin{cases} \frac{(z + z^k (\frac{q^b}{q^a}^{1-\alpha} - 1))^{1-\alpha}}{1-\alpha} & \text{if } z \geq z^k \\ \frac{z^{1-\alpha}}{1-\alpha} & \text{if } z < z^k \end{cases}$$

These preferences effectively generate a kinked relationship between x and z so that

$$z(x) = \begin{cases} q^b x^k + q^a (x - x^k) & \text{if } x \geq x^k \\ q^b x & \text{if } x < x^k \end{cases}$$

where we define $x^k = \frac{z^k}{q^b}$. This utility function tractably captures a changing tax rate. Individuals below the kink consume z at a different rate than individuals

above the kink. When $q^a < q^b$ the tax rate is declining whereas when $q^a > q^b$ the tax rate is increasing.

A number of insights follow just from writing down these preferences. First, it sheds light on why the poor might not invest in a range of high return opportunities even when they are entirely divisible (pay down their debt, use more fertilizer), and yet join ROSCAs which pay negative returns. If the preferences are such that $q^a > q^b$, so that they are more subject to temptations when they have a small amount of money at hand, than when they can spend a lot at once, then high returns on small investments are not worth much (because they will mostly go into z goods) but a ROSCA which gives them a lot of money all at once and therefore leads to a lot of x spending, is very valuable. More generally this framework can explain why there may be a preference for a certain type of illiquidity, which restricts spending unless the total amount spent is large enough. These correspond to the savings target based commitment devices studied in Ashraf, Karlan and Lin (XXXX).

Second, the popularity of micro-credit might have a lot to do with the fact it too has this structure: you get to spend a lot when you get the loan, thereby capturing a lot of x consumption and then you have to pay it back in small installments, which may just cut into z consumption. If this is the case the insistence of many MFIs on lending only for production purposes may be misdirected.

Third, the fact that the poor are more tempted, combined with the fact that more tempted save less, suggests a situation where it is very hard to get out of poverty. Indeed, in the limit it can give rise to a poverty trap. To see this, for any equilibrium, define $w_t(w)$ to be the wealth at time t in the equilibrium consumption path for an individual who is currently at wealth w . Define $\gamma(w)$ to be the current period consumption rate for an individual who has wealth w . The terms $z(w)$, $x(w)$, and $c(w)$ are defined similarly. We have left out the time subscript for these definitions since we will focus on stationary equilibria.

Our primary proposition regarding the poverty trap is:

Proposition 3 *Assume preferences such that $\alpha > 1$. Suppose $q^a < q^b$ and the parameters are such that an individual with only $q \leq q^a$ would not decumulate and an individual with a $q \geq q^b$ would not accumulate. Then for q^b large enough there exists an equilibrium with the following property: $\exists \tilde{w} > 0$, so that (1) for all $w_0 \geq \tilde{w}$, $\exists t$ so that whenever $t > t_0$ $w_t(w_0) \geq \tilde{w}$; (2) for all $w_0 < \tilde{w}$, $w_t(w_0) < \tilde{w}$ for all t .*

In words, individuals above a certain level of wealth accumulate and stay rich, whereas those below a certain level of wealth stay poor.

3.4.1 Sketch of Proof (Proof in Appendix)

Step 1: The first step is to recognize that for very general preferences, along any equilibrium accumulation path, $w_t(w)$ is monotonic—if you start richer in any period, you must end up richer as well. This follows from the fact that both u and v are concave.

Step 2: Monotonicity rules out the possibility that people go down in order to come up again. All trajectories either keep going down, always go up or end up converging to a level where they stay put.

Step 3: Defining $w^k = \frac{x^k(1+q^b)}{\gamma^b}$, which is the minimum wealth at which she will be consuming $x \geq x^k$, consider someone who starts with $w < w^k$. Suppose the globally optimal trajectory for such a person is to never cross w^k . Then the q she faces in the entire future is q^a . But then we know her global optimum will actually keep her from accumulating so that she will indeed stay under w^k .

Step 4: Now consider someone whose $w \geq w^k$. One locally optimal trajectory is for this person to choose based on the assumption that in the future the q she faces is always q^b , which, by our assumption, would indeed keep her wealth above w^k and therefore she would be right in assuming that $q = q^b$. The alternative is for her to consider another trajectory, which goes down below w^k and remains there (by step 2 she cannot climb back up). Compared to the previous alternative, this trajectory has the feature that she consumes at a higher rate than on the other trajectory, and the fraction of z consumption is higher in total consumption (here we make use of step 3). Hence the future looks worse than if she had followed the first option. In the current period she gets to consume more, but this was already an option for her when there was no risk of z consumption going up as a result and then she had not chosen it. Hence it cannot be optimal. Anyone who is above w^k will remain there.

Step 5: Now go back to someone who starts below w^k . For such a person, as we already observed one local optimum is to stay below w^k for ever. However this may not be a global optimum because if she manages to cross above w^k she will be rewarded by permanently lower z consumption. However consider a person who has a low wealth level w_0 and assume for the time being that $\theta^a = \frac{q^a}{1+q^a} = 1 - \varepsilon$, for $\varepsilon \approx 0$, so that she does not put any value on the consumption of her future selves, until the trajectory crosses w^k . If such a person had no plans of crossing w^k then her optimal choice would be to consume all of w^k . The alternative is to save until the trajectory crosses w^k . Since along such a trajectory wealth must always go up (by monotonicity), the maximum she can consume in the current period while getting on to this upwards path is given by $R(w_0 - c_0) = w_0$ or $c_0 = \frac{R-1}{R}w_0$. In other words the contemporaneous loss for someone who decides to save up is given by

$$\begin{aligned} & \frac{(1+q)}{1-\alpha} w_0^{1-\alpha} - \frac{(1+q)}{1-\alpha} \left(\frac{R-1}{R} w_0 \right)^{1-\alpha} \\ &= \frac{(1+q)}{1-\alpha} w_0^{1-\alpha} \left[1 - \left(\frac{R-1}{R} \right)^{1-\alpha} \right] = Q w_0^{1-\alpha}. \end{aligned}$$

The gain comes in at least n periods, where n is defined by $w_0 R^n = w^k$ or $n(w_0) = \frac{\log w^k - \log w_0}{\log R}$. Denoting the (finite) utility that she will get once w^k is reached by V , we can write the total gain as $\delta^{n(w_0)} V$. A necessary condition for her to be willing to save her way out of poverty is that

$$Q w_0^{1-\alpha} < \delta^{n(w_0)} V$$

Step 6: For $\alpha > 1$, $(\frac{R-1}{R})^{1-\alpha} > 1$, and therefore $Q > 0$. Therefore $Qw_0^{1-\alpha}$ goes to ∞ as w_0 goes to 0. On the other hand $n(w_0) \rightarrow \infty$, when $w_0 \rightarrow 0$, and hence $\delta^{n(w_0)}V \rightarrow 0$. Therefore the above necessary condition is violated and someone with w_0 small enough will not save their way out of poverty.

Step 7: Since the poor decumulate while the rich remain rich, by monotonicity, there must exist a critical \tilde{w} . This completes the "proof".

We can also prove a sort of a converse to this result.

Proposition 4 *Assume the same class of preferences. If $q^b > q^a$ there cannot be a poverty trap.*

This follows from the fact that it is the poor who must now accumulate faster than the rich (or decumulate slower), so the gap between them must shrink over time.

4 Conclusion: Implications of this framework

The point of this paper is to develop a simple framework that allows us to think about the role of self-control issues in understanding the behavior of the poor. Our hope is that others will find this framework as useable as we have found it, and that it can be basis for much further work on the subject.