A Dual Self Model of Impulse Control

Drew Fudenberg and David K. Levine

June 17, 2009

“The idea of self-control is paradoxical unless it is assumed that the psyche contains more than one energy system, and that these energy systems have some degree of independence from each other.”

(McIntosh [1969])

The Problem

- ♦ apparent time inconsistency that has motivated models of hyperbolic discounting

  choice between consuming some quantity today and a greater quantity tomorrow, choose lesser quantity today

  when faced with the choice between same relative quantities a year from now and a year and a day from now, choose greater quantity a year and a day from now.

- ♦ Rabin's [2000] paradox of risk aversion in the large and small

  the risk aversion experimental subjects show to very small gambles implies hugely unrealistic willingness to reject large but favorable gambles

Overview

- view decision problems as a game between a sequence of short-run impulsive selves and a long-run patient self who controls at a cost the short-run self's preferences
- consistent with MRI evidence
- similar to many recent models
- consistent with Gul-Pesendorfer axioms
- benefit of commitment – current short-run self does not care about a year versus a year and a day, so no cost to long-run self of committing
- but short-run self does care about today but not tomorrow, so costly to get the short-run self to forgo consumption today in exchange for consumption tomorrow

The Model

time discrete and unbounded, $t = 1, 2, \ldots$.

fixed, time-and history invariant set of actions A for the short-run selves

a measure space Y of states

a set R of self-control actions for the long-run self, $0 \in R$ means no self-control is used

A, Y, R closed subsets of Euclidean space

finite history of play $h \in H$ of the past states and actions, $h = (y_1, a_1, r_1, \ldots, y_t, a_t, r_t)$ plus the null history 0

H_t the set of t-length histories H_t

length of the history $t(h)$, final state in h is $y(h)$, initial state y_1

probability distribution over states at $t+1$ depends on time-*t* state and action y_t, a_t by stochastic kernel $\mu(y, a)$

note that the long-run self's action r has no effect on states game is between long-run self with strategies $\sigma_{LR} : H \times Y \rightarrow R$

and sequence of short-run selves

period t short-run self plays in only one period, observes self-control action of long-run self prior to moving; uses strategy $\sigma_t : H_t \times Y \times R \rightarrow A$

collection of one for each SR is denoted σ_{SR}

for every measurable subset $R' \subseteq \mathbf{R}, A' \subseteq \mathbf{A}$ the functions $\sigma_{LR}(\cdot,\cdot)[A'], \sigma_t(\cdot,\cdot,\cdot)[R']$ are measurable

strategies together with measure μ give rise to a measure π_t over length t histories utility of the short-run self is $u(y,r,a)$: long-run player's self-control action influences the short-run player's payoff

$u_t(h) = u(y(h), \sigma_{LR}(h, y(h)), \sigma_t(h, y(h), \sigma_{LR}(h, y(h)))$

utility of the long-run self is

$$U_{LR}(\sigma_{LR}, \sigma_{SR}) = \sum_{t=1}^{\infty} \delta^{t-1} \int u_t(h) d\pi_t(h)$$

no intrinsic conflict between long-run and short-run self

Assumption 0 (Upper Bound on Utility Growth): For all initial conditions

$$\sum_{t=1}^{\infty} \delta^{t-1} \int \max\{0, u\ (h)\} d\pi_t(h) < \infty.$$

short-run self optimizes following every history: *SR-perfect*

interested in SR-perfect Nash equilibria

Assumption 1 (Costly Self-Control): If $r \neq 0$ then $u(y, r, a) < u(y, 0, a)$.

Assumption 2 (Unlimited Self-Control): For all y, a there exists r such that for all a', $u(y, r, a) \geq u(y, r, a')$.

with these two assumptions we may define the cost of self-control

$$C(y, a) \equiv u(y, 0, a) - \sup_{\{r | u(y,r,a) \geq u(y,r,\cdot)\}} u(y, r, a)$$

Assumption 3 (Continuity): $u(y, r, a)$ is continuous in r, a.

the supremum can be replaced with a maximum Assumptions 1 & 3 imply cost continuous and

Property 1: (Strict Cost of Self-Control) If $a \in \arg\max_{a'}(u(y, 0, a'))$ then $C(y, a) = 0$, and $C(y, a) > 0$ for $a \notin \arg\max_{a'}(u(y, 0, a'))$.

Assumption 4 (Limited Indifference): for all $a' \neq a$, if $u(y,r,a) \geq u(y,r,a')$ then there exists a sequence $r^n \rightarrow r$ such that $u(y,r^n,a) > u(y,r^n,a')$.

short-run self is indifferent, long-run self can break tie for negligible cost *reduced-form* optimization problem

$H^{AY} = \{(y_1, a_1, \ldots, y_t, a_t)\}_t$ *reduced* histories

problem of choosing a strategy from reduced histories and states to actions, $\sigma_{RF} : H^{AY} \times Y \to \mathbf{A}$, to maximize the objective function

$$U_{RF}(\sigma_{RF}) = \sum_{t=1}^{\infty} \delta^{t-1} \int [\, u(y(h), 0, a) - C(y(h), a)\,] d\sigma_{RF}(h, y(h))[\, a\,] d\pi_t(h)$$

Theorem 1 (Equivalence of Subgame Perfection to the Reduced Form): Under Assumptions 1-4, every SR-perfect Nash equilibrium profile is equivalent to a solution to the reduced form optimization problem and conversely.

Assumption 5 (Opportunity Based Cost of Self Control) If

$\max_{a'} u(y,0,a') \geq \max_{a'} u(y',0,a')$ and $u(y,0,a) \leq u(y',0,a)$ then $C(y,a) \geq C(y',a)$.

This assumption says that the cost of self control depends only on the utility of the best foregone utility and the utility of the option chosen

Adding Assumption 5 to Assumptions 1-3 implies a continuous function $C(y,a) = \tilde{C}(u(y,0,a), \max_{a'} u(y,0,a'))$

decreasing in realized utility, increasing in temptation, $\tilde{C}(u,u) = 0$

Assumption 5 (Linear Self-Control Cost):
$C(y,a) = \gamma\left[\max_{a'} u(y,0,a') - u(y,0,a)\right]$

Reduced Form of the Model

Summary:

Let y be that state and a be the action taken at that state. Under various assumptions the game between the short-run and long-run self is reducible to an optimization problem with control cost for the long-run self

$$
\begin{aligned}
U &= \sum_{t=1}^{\infty} \delta^{t-1} \int \left[u(y,0,a) - C(y,a) \right] d\pi_t(y(h)) \\
&= \sum_{t=1}^{\infty} \delta^{t-1} \int \left[(1+\gamma) u(y,0,a) - \gamma \max_{a'} u(y,0,a')) \right] d\pi_t(y(h))
\end{aligned}
$$

A Simple Banking Model and The Rabin Paradox

many ways of restraining short-run self besides the use of self-control

make sure the short-run self does not have access to resources that would represent a temptation

The Environment

each period consists of two subperiods: "bank" subperiod and "nightclub" subperiod

during "bank" subperiod

- consumption is not possible
- wealth y_t is divided between savings s_t, which remains in the bank, and "pocket" cash x_t which is carried to the nightclub

at the nightclub

- consumption $0 \leq c_t \leq x_t$ is determined, with $x_t - c_t$ returned to the bank at the end of the period
- wealth next period is just $y_{t+1} = R(s_t + x_t - c_t)$

- discount factor between two consecutive nightclub is δ
- preferences are logarithmic

perfect foresight problem savings only source of income

- no consumption possible at bank
- long-run self gets to call the shots
- can implement a^*, the optimum of the problem without self-control, simply by choosing pocket cash $x_t = (1 - a^*)y_t$ to be the target consumption
- it is the case that the short-run self will in fact spend all the pocket cash; that having solved the optimum without self-control, the long-run self does not in fact wish to exert self-control at the nightclub.

stochastic cash receipts (or losses)

at the nightclub in the first period there a small probability the agent will be offered a choice between several lotteries

$\tilde{z}_1$ be the chosen lottery

[if choices are drawn in an i.i.d. fashion, results in a stationary savings rate tslightly different from the a^*above; if probability that a non-trivial choice is drawn is small, savings rate will be very close to a^*]

consider the limit where the probability of drawing the gamble is zero; avoid an elaborate computation to find a savings rate close to but not exactly equal to a^*.

behavior conditional on each possible realization z_1

short-run self constrained to consume $c_1 \leq x_1 + z_1$

first order condition for optimal consumption gives

$$c_1 = \left(1 - \frac{\delta}{\delta + (1+\gamma)(1-\delta)}\right)(y_1 + z_1) \equiv (1 - B)(y_1 + z_1)$$

if c_1 satisfies the constraint $c_1 \leq x_1 + z_1$ it represents the optimum; otherwise the optimum is to consume all pocket cash, $c_1 = x_1 + z_1$

$c_1 \leq x_1 + z_1$ if $z_1 \geq z_1^*$, where the critical value of z_1^* is

$$z_1^* = \gamma(1-\delta)y_1$$

Theorem 2: If $z_1 < z_1^*$, overall utility is

$$\log(x_1 + z_1) + \frac{\delta}{(1-\delta)}\left(\log(1-\delta) + \log(R(y_1 - x_1)) + \frac{\delta}{1-\delta}\log(R\delta)\right) \quad \textbf{(6)}$$

If $z_1 > z^*$ utility is

$$\begin{aligned}&(1+\gamma)\log\left(\frac{(1-\delta)(1-\gamma)}{1+\gamma(1-\delta)}(y_1 + z_1)\right) - \gamma\log(x_1 + z_1)\\&+\frac{\delta}{(1-\delta)}\left(\log(1-\delta) + \log\left(\frac{R\delta}{1+\gamma(1-\delta)}(y_1 + z_1)\right) + \frac{\delta}{1-\delta}\log(R\delta)\right)\end{aligned} \quad (7)$$

risk aversion

$\tilde{z}_1 = \overline{z} + \sigma\varepsilon_1$,

ε_1 has zero mean and unit variance, σ is very small

comparing a lottery with certainty equivalent

For $\overline{z} < z^*$ overall payoff is given by (6)

relative risk aversion constant and equal to ρ

wealth is $w = x_1 + \overline{z}_1$ so risk is measured relative to pocket cash for $\overline{z} > z^*$, the utility function (7) is the difference between two utilitity functions, one of which exhibits constant relative risk aversion relative to wealth $y_1 + \overline{z}$, the other of which exhibits constant risk aversion relative to pocket cash $x_1 + \overline{z}$

γ is small, the former dominates, and to a good approximation for large gambles risk aversion is relative to wealth, while for small gambles it is relative to pocket cash

Rabin [2000]

"Suppose we knew a risk-averse person turns down 50-50 lose \$100/gain \$105 bets for any lifetime wealth level less than \$350,000, but knew nothing about the degree of her risk aversion for wealth levels above \$350,000. Then we know that from an initial wealth level of \$340,000 the person will turn down a 50-50 bet of losing \$4,000 and gaining \$635,670."

The point being of course that many people will turn down the small bet, but no one would turn down the second. In our model, however, we can easily explain these facts, with, say, logarithmic utility.

small stakes gamble

- first bet isensibly interpreted as a pocket cash gamble
- experiments with real monetary choices in which subjects exhibit similar degrees of risk aversion over similar stakes are
- if the agent not carrying $100 in cash, transaction cost in the loss state of finding a cash machine or bank
- easiest calculations are when gain $105 is smaller than threshold z^*
- logarithmic utility requires the rejection of the gamble if pocket cash x_1 is $2100 or less
- for gain of $105 is to be smaller than the threshold z^*, $\gamma \geq 105 / x_1$

- for pocket cash $x_1 = 2100$ need $\gamma > .05$
- for pocket cash equal to daily atm withdrawal limit $x_1 = 300$, need γ at least 0.35
- calculations quite robust to the presence of the threshold
- for pocket cash is \$300, wealth \$300,000 and $\gamma = 0.05$ then favorable state of \$105 well over the threshold of \$15
- computation shows that the gamble should still be rejected
- not even close to the margin

large stakes gamble

- unless pocket cash at least $4,000 second gamble must be for bank cash
- for bank cash relevant parameter wealth, not pocket cash
- if wealth is at least $4,026 second gamble will always be accepted
- for example, an individual with pocket cash of $2100, $\gamma = 0.05$ and wealth of more than $4,026 will reject the small gamble and take the large one
- for example, an individual with pocket cash of $300, $\gamma = 0.05$ and wealth equal to the rather more plausible $300,000 will also reject the small gamble and take the large one