

A Dual Self Model of Impulse Control

Drew Fudenberg and David K. Levine

September 30, 2005

“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, $t = 1, 2, \dots$

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, $h = (y_1, a_1, r_1, \dots, y_t, a_t, r_t)$ 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

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

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

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

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 (slightly 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

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

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

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

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

comparing a lottery with certainty equivalent

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

relative risk aversion constant and equal to ρ

wealth is $w = x_1 + \bar{z}_1$ so risk is measured relative to pocket cash

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

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 is sensibly 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

Discussion

- ◆ assumed cash only available at banking stage
- ◆ if agent, when banking, anticipates the availability of \$300 from an ATM during the nightclub stage, it is optimal to reduce pocket cash by this amount
- ◆ if the goal is to have pocket cash less than \$300, then self-restraint will be necessary in the presence of cash machines
- ◆ which is why we find cash machines where impulse purchases are possible
- ◆ in equilibrium, few if any, additional overall sales are induced by the presence of these machines, since their presence is anticipated, but the competitor who fails to have one will have few sales
- ◆ one consequence of the dual self-model is that we may see an inefficiently great number of cash machines.

- ◆ credit cards and checks pose complications
- ◆ for many people future consequences of using credit cards and checks significantly different than expenditure of cash
- ◆ it is one thing to withdraw the usual amount of money from the bank, spend it all on the nightclub and skip lunch the next day
- ◆ something else to use a credit card at the nightclub, which, in addition to the reduction of utility from lower future consumption, may result also in angry future recriminations with one's spouse, or in the case of college students, with the parents who pay the credit card bills
- ◆ so often optimal to exercise a greater degree of self-control with respect to non-anonymous expenditures such as checks and credit cards, than with anonymous expenditures such as cash
- ◆ consistent with the finding of Wertenbroch, Soman, and Nunes [2002] that individuals who are purchasing a good for immediate enjoyment have a greater propensity to pay by cash, check or debit card than by credit card

Procrastination

every period $t = 1, 2, \dots$ short-run self must either take an action (“write the great new novel”) or wait

waiting allows the self to enjoy a leisure activity yielding a stochastic amount of utility x_t , whose value is known at the start of that period

for example the leisure activity is playing outside and utility depends on the weather

the x_t are i.i.d. with fixed and known cumulative distribution function P and associated density p on the interval $[\underline{x}, \bar{x}]$

taking the action ends the game, and gives a flow of utility v beginning next period

corresponding present value $\frac{\delta}{1 - \delta} v \equiv \delta V$.

waiting causes the problem repeats in the next period

current value of x has a monotone effect on the payoff to waiting, no effect on the payoff to doing it now, so optimal solution is a cutoff rule:

$x \geq x^*$ wait, $x < x^*$ take the action

maximum utility in any period is x_t , the payoff to waiting; doing it now requires foregoing x_t . So waiting has no self-control cost, and acting has self-control cost of γx_t

Theorem 3 (i) $\underline{x} \leq x^* < \bar{x}$; if $\delta v > \delta \mu + (1 - \delta)(1 + \gamma)\underline{x}$ then $x^* > \underline{x}$.

(ii) When $\underline{x} < x^* < \bar{x}$, $\frac{dx^*}{d\gamma} < 0$, so expected waiting time is increasing in the cost of self control.

Miao [2004] applies a dual-self model to a very similar problem in which the reward is stochastic and the cost fixed (the “search” model)

O'Donoghue and Rabin [2001] analyze implications of hyperbolic discounting in a similar stopping time problem

they say the agent “procrastinates” if he never acts even though there is an action that is worth doing *given* his hyperbolic discounting of future returns

their model has multiple equilibria; equilibria are cyclic, with intervals of length T between “action dates” these cyclic equilibria seem artificial and unappealing

they restrict attention to equilibria that are limits of equilibria in the finite horizon

restriction relies on long chains of backwards induction and is not robust

despite the presence of multiple equilibria, O'Donoghue and Rabin can show that sophisticated agents (that is those who know their own hyperbolic parameter β) never procrastinate, although they may postpone acting for a few periods

DellaVigna and Malmendier [2003] calibrate the O'Donoghue-Rabin model to data on delay in canceling health club memberships

they use hyperbolic preferences and “lack of sophistication,” meaning that consumers misperceive their own hyperbolic parameter and thus incorrectly forecast their health club usage

our model suggests several qualifications to their analysis

1. as is standard in models of timing, it is not in general optimal for the agent to act whenever he is indifferent between acting now or not at all, as there is an “option value” in waiting.
2. there is evidence that agents do not have perfect knowledge about themselves. We do expect them to have more information about things that they have had more chances to observe. So is it reasonable to suppose that misperceptions about impulsiveness are more likely than misperceptions about the short-run disutility and long-run benefits of going to the health club?

Cognitive Load and Self Control

Shiv and Fedorikhin [1999]

subjects asked to memorize either a two- or a seven-digit number, and then walk to a table with a samples of two deserts, namely chocolate cake and fruit salad

subjects then pick a ticket for one of the deserts, and go to report both the number and their choice in a second room

subjects who were asked to remember the seven-digit number chose cake 63% of the time, while subjects given the two-digit number chose cake 41% of the time, and this difference was statistically significant

possible actions be h (chocolate) and f (fruit)

short-term utilities $u^h > u^f$

assume long-term utility of fruit is higher than chocolate

state y is cognitive load d

assumption 5 self-control cost of f with load d

$$C(d, f) = c(d, u^h, u^f)$$

one explanation is

$$C(d, f) = d \cdot |u^h - u^f|$$

unsatisfactory

cognitive resources for memorization increases the marginal cost of self-control

expect that using these same resources for self-control should also change the marginal cost of self control

prefer the non-linear specification

$$C(d, f) = g(d + u^h - u^f) - g(d),$$

where g is an increasing convex function with $g' > 0$ and $g'' > 0$.

violates linearity

Gul-Pesendorfer Axioms

initial period action a has no utility consequences for the short-run self just determines state y and utility possibilities starting in period 2

for any initial choice consider a plan of action starting in period 2

results in a second-period utility to the short-run self of u , and a present value of V to the long-run self starting in period 3

long-run agent cares about feasible utility consequences (u, V)

think of initial period choice as choosing a set of feasible (u, V) pairs

our model: utility to the initial long-run self choosing the set W is

$$\Phi(W) = \max_{(u, V) \in W} [u + \delta V - C | \max(u' | (u', V') \in W | , u)].$$

gives rise to a preference ordering over sets W

Gul and Pesendorfer [2001] show that their axioms equivalent to a representation

$$\max_{(u,V) \in W} [h(u,V) + H(u,V)] - \max_{(u,V) \in W} H(u,V).$$

When

$$C \mid \max(u' \mid (u',V') \in W), u = \gamma \mid \max(u' \mid (u',V') \in W) - u),$$

we take $h(u,V) = u + \delta V$ and $H(u,V) = \gamma u$

Non-linear C in our model is consistent with neither their axioms nor weaker ones of Dekel, Lipman, and Rustichini

Set-Betweenness

Theorem 5: Under Assumptions 1-5, the induced preferences Φ over choice sets satisfy set-betweenness. That is, for all choice sets W, Z either $W \succeq W \cup Z \succeq Z$ or $Z \succeq W \cup Z \succeq W$.

Just copy their intuitive proof.

Shows that Assumption 5 rules out the preferences in Dekel, Lipman, and Rustichini's Example 1,

$$\{b\} \succ \{b, h\}, \{b\} \succ \{b, p\}$$

$$\{b, p\} \succ \{b, h, p\}$$

$$\{b, h\} \succ \{b, h, p\}$$

long-run self is uncertain which of h and p will be more tempting

our model requires that no uncertainty realized between long-run self choice r and short-run choice a

Independence Axiom

violation of the independence axiom on sets of lotteries used by Gul and Pesendorfer and Dekel, Lipman, and Rustichin

a lottery that takes place after a chosen by the short-run self

u, V represent random variables

what matters to either the short- or long-run self is expected present value of this lottery

$$\Phi(W) = \max_{(u,V) \in W} | u + \delta V - C | \max(Eu' | (u',V') \in W | , Eu) |.$$

a non-linear function of an expected value so violates independence

closely connected to Machina [1984]

shows ranking over lotteries induced *ex ante* before a decision generally violates the independence axiom even though independence axiom is satisfied *ex post*

dessert options replaced by lotteries with probability p of the chosen dessert and probability $1 - p$ of no dessert

our model with convex costs implies that more agents chose fruit when probability p of a dessert is lower – the self-control problem is mitigated by the uncertainty

Independence of Irrelevant Alternatives

Example 2 of Dekel, Lipman, and Rustichini

three possible actions, broccoli (b), frozen yogurt (y), and ice cream (i),
with $\{b, y\} \succ \{y\}$ and $\{b, i, y\} \succ \{b, i\}$

frozen yogurt a “compromise” option appealing in the face of strong temptations but not weaker ones

Dekel, Lipman, and Rustichini show not consistent with the Gul and Pesendorfer axioms

consistent with our Assumptions 1-5

$$(u, V)(b) = (0, 100); (u, V)(y) = (8, 30); (u, V)(i) = (14, 0)$$

$$\Phi(W) = \max_{(u, V) \in W} [u + .9V - .5 | \max(u' | (u', V') \in W) - u]^2]$$

$\{b\}$ vs $\{y\}$ Ir value of broccoli is $90 - .5(64) = 58$

value of yogurt 35 in both $\{b, y\}$ and $\{y\}$

in $\{b, y, i\}$ value of yogurt $35 - .5(36) = 17$; value of ice cream 14; value of broccoli $90 - .5(196) = -8$.

Our choice function violates the independence of irrelevant alternatives

choose y from $\{b, i, y\}$ but b from $\{b, y\}$

Dekel, Lipman and Rustichini propose different explanation based on the idea of that long-run self uncertain of the short-run self's choice function

their explanation satisfies form of independence of irrelevant alternatives for stochastic choice functions

when self-control costs are linear action chosen from set W is $\arg \max_{(u,V) \in W} [(1 + \gamma)u + \delta V - \gamma \max(u' | (u', V') \in W)]$

temptation $\max(u' | (u', V') \in W)$ is “sunk cost” ; does not effect plan (u, V) chosen from W

linear cost has independence of irrelevant alternatives satisfied.