Behavioural Economics **University of Oxford**
Vincent P. Crawford **Michaelmas Term 2012**

Behavioural Decision Theory: Present-Bias and Time-Inconsistency in Intertemporal Choice
(with very large debts to David Laibson and Ted O'Donoghue; material from their slides is used with permission)

Revised 18 October 2013

The standard model of intertemporal choice in economics is time-separable utility with exponential discounting (Ramsey 1928 *EJ*, Samuelson, 1937 *REStud*):

$$U_t = u_t + \delta u_{t+1} + \delta^2 u_{t+2} + \delta^3 u_{t+3} + \cdots$$

With exponential discounting, your trade-offs between receiving utils today and receiving them with delay are independent of when that delay occurs:

- If you prefer 10 utils now to 15 next week, then you must also prefer 10 utils in 25 weeks to 15 in 26 weeks.

This agreement implies "time-consistent" (or "dynamically consistent") choice, in that early and later selves agree on the ranking of plans, so that the mere passage of time (with no new information) doesn't change the ranking.

(In more general models, receiving new information over time could affect your preferences over plans conditional on the information received.

But if your preferences are time-consistent, receiving new information over time cannot alter your preferences over plans contingent on resolution of uncertainty.)

Thus, an exponential planner can simply maximize lifetime utility at the start of his life without worrying about later selves overturning his decisions.

For an exponential discounter, if it's beneficial to do something next week/month/year/etc., other things equal it's (even more) beneficial to do it Now.

But people often seem to deviate systematically from these implications of the exponential discounting model.

Moe:

"This thing can **flash fry** a **Water Buffalo** in **40 seconds.**"

Homer:

"Ohhhhh, **40 seconds**!
But I want mine Now."

There are many familiar quotations you would never hear from an exponential discounter (this intuitive “evidence” will later be supplemented by real evidence):

- Next month, I’ll quit smoking.
- Next week, I’ll catch up on the required reading.
- Tomorrow morning, I’ll wake up early and exercise.
- After Christmas, I’ll start eating better.
- Next weekend, I’ll send in this rebate form.
- Next month, I’ll start saving for retirement.

These quotations all involve a pattern of deviations from time-consistent exponential discounting called “present bias”:

• Excessively favoring gratification (or avoiding nongratification) now at the expense of future gratification.

“Excessive” means “relative to exponential”: Exponential discounters do prefer present to future utils, other things equal; just not as much as present-biased people.

The bias is *systematically* (if not universally) in favor of the present.

Thus it can't usefully be treated as random deviations from an exponential model.

Future bias would yield quotations you seldom (never?) hear:

- I plan to watch more TV next year.
- I plan to eat more cookies and doughnuts next year.
- I plan to smoke more cigarettes next year.
- I plan to borrow more on my credit card next year.
- I plan to exercise less next year.
- I plan to wake up later next year.

One possible exception:

- John Maynard Keynes' last words: "I wish I had drunk more champagne"…but he was a workaholic, so present bias might have put work before champagne.

The quotations you actually hear also reflect a time-*in*consistent tension between the preferences of current and future selves:

- Early selves plan to further long-term goals but later selves may sacrifice those goals in favor of instant gratification.

- Present bias often leads people to put off unpleasant tasks that yield future benefits, a deviation from time-consistent planning called procrastination.

- E.g. “Da mihi castitatem et continentiam, sed noli modo.” (“Give me chastity and continence—but not yet.”)—Saint Augustine of Hippo

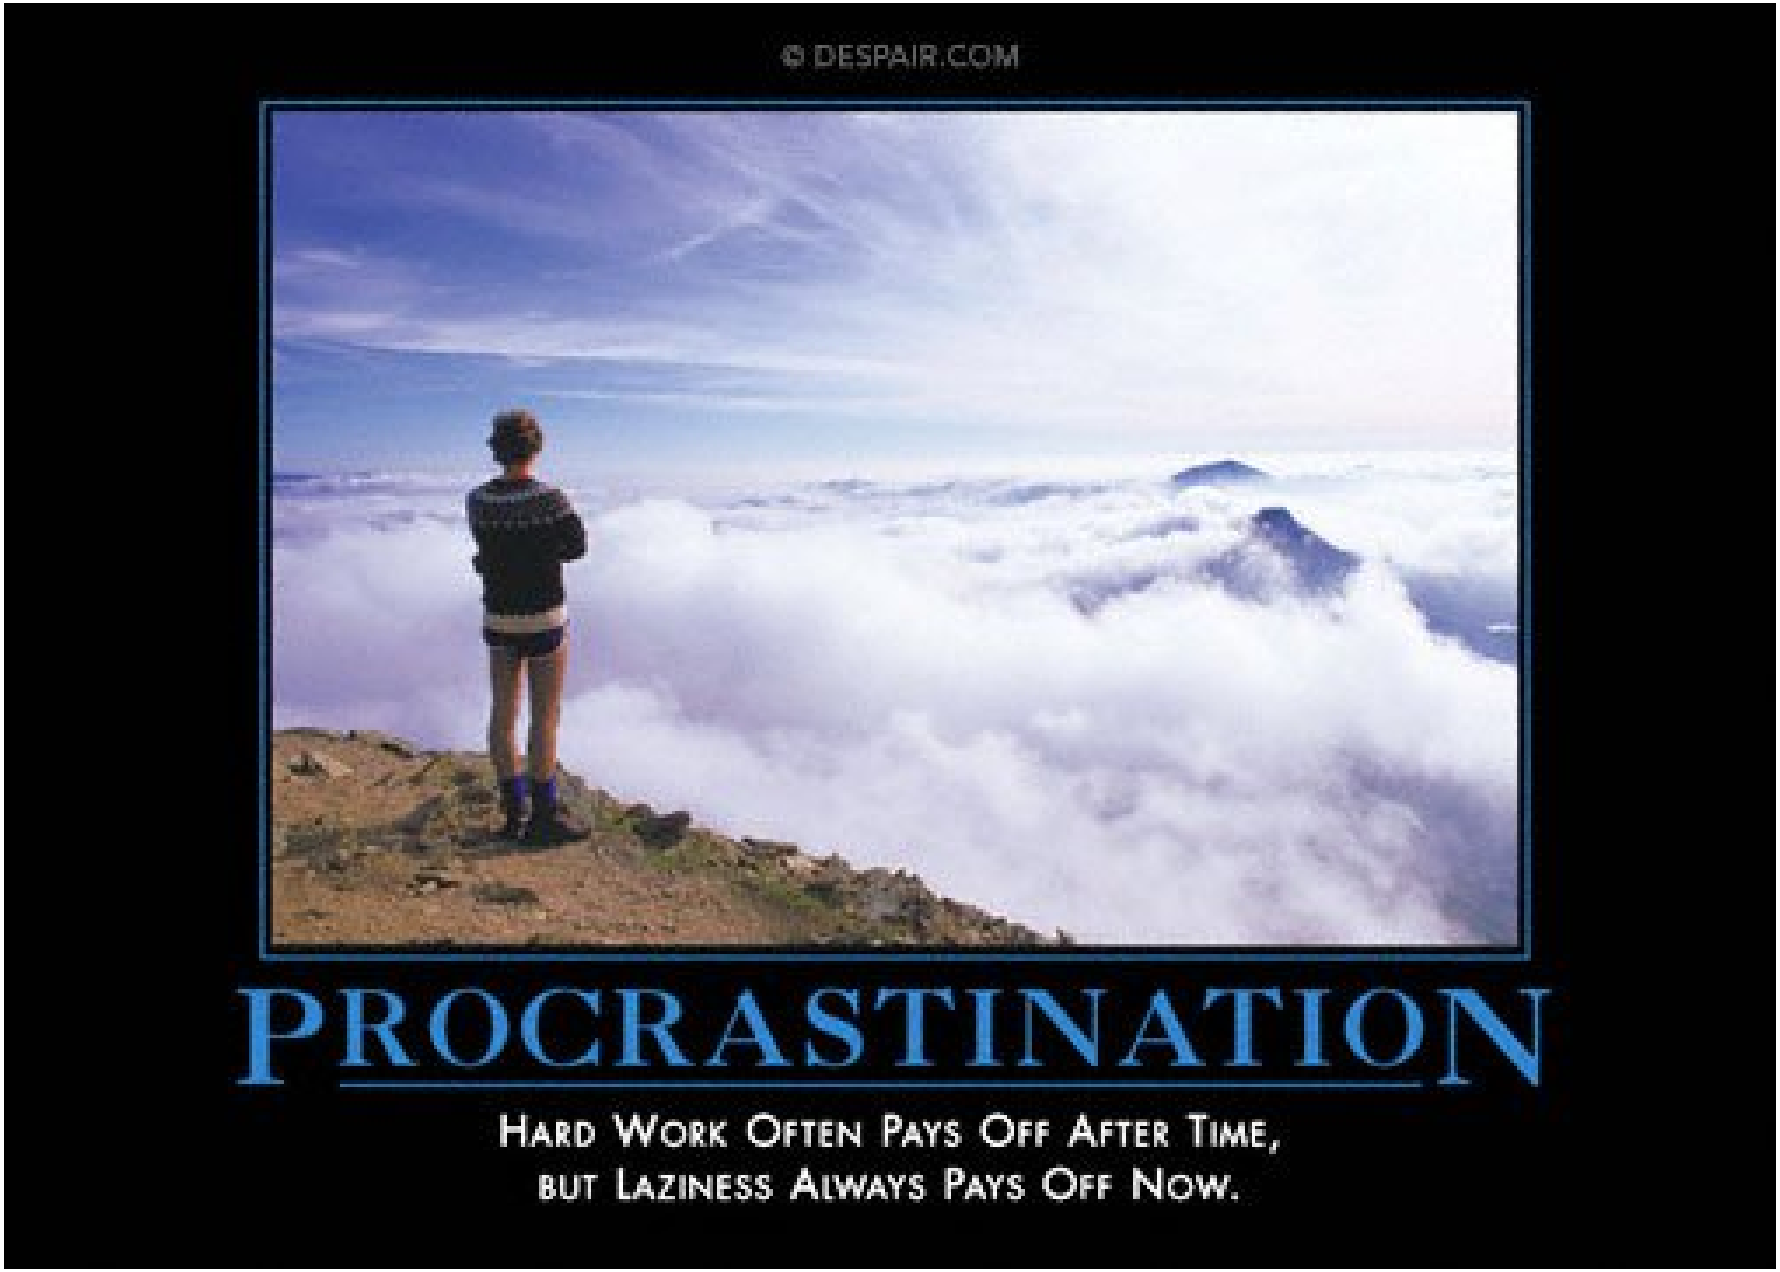
© DESPAIR.COM
PROCRASTINATION
HARD WORK OFTEN PAYS OFF AFTER TIME,
BUT LAZINESS ALWAYS PAYS OFF NOW.

Procrastination is an important problem in modern life (numbers from 2011):

- Among common students' problems, procrastination, with 14,800,000 Google hits, is 33% more important than ennui with 11,100,000, but 47% less important than plagiarism with 27,800,000.

- More generally procrastination is far less important than pride with 418,000,000; lust 216,000,000; wrath 148,000,000; greed 114,000,000; or envy 110,000,000.

- It runs close behind sloth with 15,900,000 (not counting the sloths with toes), but easily tops gluttony with 5,450,000.

- (For calibration: Obama 723,000,000 Google hits; Lady Gaga 479,000,000; The Simpsons 159,000,000; Emma Watson 92,300,000; Marx (not counting Groucho) 90,900,000; Keynes 51,000,800.)

Future bias would make people accelerate unpleasant tasks or defer gratification.

English has a word for this, but possibly you haven't heard it yet:

- "Preproperation" (456 Google hits and counting).

Discount functions and discount rates

- Discount function: *u* utils in τ periods are worth *D*(τ) utils today.

- The discount rate is the rate of decline in the discount function, the rate at which the value of a util declines with delay:

$$\rho(\tau) \equiv -\frac{dD(\tau)/d\tau}{D(\tau)}$$

- Discounting is a feature of preferences, not intertemporal transformation possibilities, which are determined by the interest rate.

- With exponential discounting $D(\tau) \equiv \delta^{\tau}$, for which, when $\delta \simeq 1$,

$$\rho(\tau) \equiv -\frac{dD(\tau)/d\tau}{D(\tau)} \equiv -ln\delta \simeq 1-\delta,$$

so the discount rate does not change with the horizon.

- If discounting is stationary (i.e. if we only care about delay from "now", whenever "now" is), then exponential discounting is the only discount function that yields time-consistent preferences (Strotz, 1956 *REStud*).

Can real people's discounting actually be exponential?

Many people seem to have at least a 1% preference for gratification today over gratification tomorrow.

But with exponential discounting, this would imply a yearly discount factor of $0.99^{365} \simeq 0.026$.

- So 100 utils in a year would be worth 2.6 utils today.

- So 100 utils in 10 years would be worth 1×10^{-14} utils today.

Even $0.999^{365} \simeq 0.694$; and $0.99^{52} \simeq 0.593$.

With exponential discounting, plausible discounting for a year from now implies implausibly low (virtually zero) discounting for a day or a week from now.

Experimental evidence

Thaler (1981 *Economics Letters*) asked subjects to choose between money now versus more money later hypothetically (but many similar experiments have now used real rewards, with similar results):

- What amount makes you indifferent between \$15 today and \$*X* in 1 month?

  Typical response: $X = 20$.

The baseline exponential discounting model has $Y = \delta^t X$. Recall that to relate a discount factor $\delta = e^{-\rho}$ to a discount rate ρ, $\ln \delta = \ln e^{-\rho} = -\rho$.

So for $X = 20$, $\rho = -\ln \delta = \frac{1}{t} \ln \frac{X}{15} = 12 \ln \frac{20}{15} \simeq 345\%$ per year.

- What amount makes you indifferent between \$15 today and \$*X* in ten years?

  Typical response: $X = 100$. Implied discount rate $\simeq 19\%$ per year.

But “money now vs. money later” has many confounds:

- If an agent has access to a perfect capital market, he should maximize net present value and then trade to the intertemporal optimum; but present value reflects market interest rates, not time preferences. (In practice, dynamic choices involving money don’t even seem to measure the interest rate.)

- Unreliability of future rewards (trust)

- Curvature or nonseparability of utility function (Andreoni and Sprenger 2012 *AER*)

Dynamic choices in kind (ice cream at t or more at t + τ) are also confounded via trust and utility function curvature, but probably better than money experiments.

McClure, Ericson, Laibson, Loewenstein and Cohen (2007 *Journal of Neuroscience*) studied very thirsty subjects:

- 60% of subjects chose juice now over twice the juice in 5 minutes.

- 30% of subjects chose juice in 20 minutes over twice the juice in 25 minutes.

Estimated 5-minute discount rate is 50% and the “long-run” discount rate is 0%.

Inferences from static choices in kind (e.g. chocolate at t or orange at t) are probably still more reliable.

Read and van Leeuwen (1998 *Organizational Behavior and Human Decision Processes*):

- Overwhelming majority of subjects chose a healthy snack now to eat next week.

- Overwhelming majority of subjects chose an unhealthy snack now to eat now.

Badger et al. (2007 *J. Health Econ.*) on recovering heroin addicts' monetary equivalent of an extra dose of Bupronorphine (heroin substitute like Methadone, reduces craving, aiding withdrawal).

Elicited willingness to pay for second dose from 13 long-time heroin addicts regularly receiving single dose (but this 2nd dose is still attractive to addicts): 2nd dose for $10? $20? ... $100?

Subjects told (truthfully) that one of their choices, randomly selected, would be implemented; hence, had incentives to choose according to true preferences.

Half asked when more deprived (2 hours before scheduled dose), half when less deprived (right after scheduled dose). Half asked for 2nd dose today, half for 2nd dose on next visit.

Always asked about second dose, always delivered in satiated state (exact same circumstances, by experienced addicts).

	Immediate	Delayed (+5)
Satiated	$50	$35
Deprived	$75	$60

Present bias. Also projection bias. What is the true value?

Probably the $35. Not because drugs are bad and this is lowest, but because present bias says 2nd column better than first, and projection bias says bottom row better than top.

Deciding in same emotional/craving state as situation, but ahead of time so self-control not an issue, is often the ideal.

Note limitations of revealed preference.

Read, Loewenstein & Kalyanaraman's (1999 J. *Behavioral Decision Making*) subjects chose among 24 movie videos.

- Some were "lowbrow": e.g *Four Weddings and a Funeral*
- Some were "highbrow": e.g. *Schindler's List*

Canonical highbrow movie:

(A man seeks answers about life, death, and the existence of God as he plays chess against the Grim Reaper during the Black Death.)

Results:

- Choosing for tonight: 66% of subjects chose lowbrow.

- Choosing for next Tuesday: 37% chose lowbrow.

- Choosing for second Tuesday: 29% chose lowbrow.

Evidence on time-inconsistency, demand for commitment, and naivete about behaviour of one's own future selves

Present bias has different implications when people are aware of their own bias than when they are not.

Also, to the extent that people are aware of their bias and the time-inconsistency it causes, it should create a demand for commitment, which if feasible is always valuable from the standpoint of the current self.

Dellavigna and Malmendier (2006 *AER*) found, analyzing a field panel data set that tracks health club members' usage over time:

- Average cost of gym membership: $75 per month
- Average number of monthly visits: 4
- Average cost per vist: $19
- Cost of "pay per visit": $10

The monthly contract has automatic renewals, and there seems to be procrastination in cancelling in the form of a lag between last usage and cancellation that is positively correlated with overpayment in the initial months.

Present bias is not enough to explain these results by itself, because if you correctly predicted your future decisions, you wouldn't join in the first place.

Also need partial naivete about behavior of future selves, as discussed below.

Ariely and Wertenbroch (2002 *Psychological Science*) ran experiments using course term papers and proofreading tasks.

Term paper subjects were 99 professionals in an executive-education course at MIT, where three short papers were required.

Each paper had a deadline, with a 1% grade penalty per day late for all subjects.

Two treatments were run between subjects in different sections:

- No Choice: Exogenously imposed, evenly spaced deadlines.

- Free Choice: Each student chose his own deadlines.

Results

In the free-choice group, 37/51 people imposed deadlines on themselves.

- Average deadline for paper 1 was 42 days before end of term on average.
- Average deadline for paper 2 was about 26 days before end of term.
- Average deadline for paper 3 was about 10 days before end of term.

On average the grades in the no-choice section were 89, versus 86 in the free-choice section.

- People chose to make costly commitments, which is consistent with present bias and some degree of sophistication.
- But their chosen commitments were far from optimal, suggesting some naivete.

Subjects were also recruited for the proofreading tasks and paid for performance, with a $1 penalty per day late for all subjects, and a $0.10 penalty for errors:

> "Sexual identity is intrinsically impossible," says Foucault; however, according to de Selby [1], it is not so much sexual identity that is intrinsically impossible, but rather the dialectic, and some would say the satsis, of sexual identity. Thus, D'Erlette [2] holds that we have to choose between premodern dialectic theory and subcultural feminism imputing the role of the observor as poet."

There were three such texts, and three treatments were run between subjects in different groups:

- No Choice: Exogenously imposed, evenly spaced deadlines of seven days for each text.
- End deadline: All three texts due at the end of 21 days.
- Free Choice: Each student chose his own deadlines for each text, within the 21-day window.

Results

- Subjects in the free-choice group spaced out their deadlines.

- Performance (both freedom from errors and on-time delivery) was highest in the no-choice treatment, followed by the free choice and end deadline treatments.

- Again subjects chose to make costly commitments, but their chosen commitments were far from optimal.

The results of both studies suggest that subjects had some present bias and some degree of sophistication, but also some naivete about their own behavior.

Quasi-hyperbolic discounting

What kind of model can explain these patterns?

Consider a model with "quasi-hyperbolic" discounting (Phelps and Pollak 1968 *REStud*, Laibson 1997 *QJE*, O'Donoghue and Rabin 1999 *AER*; a.k.a. "hyperbolic" or "β,δ"):

$$D(\tau) \equiv \begin{cases} 1 \text{ if } \tau = 0 \\ \beta\delta^{\tau} \text{if } \tau = 1,2, \ldots \end{cases}$$

(Tractably approximates (non-quasi-)hyperbolic discounting : D(x) = 1/(1 + kτ).)

Plausibly, β << 1 and δ < 1. E.g. β = 2/3 (for one day) and δ ≃ 0.95.

• We can write a quasi-hyperbolic utility function as

$$U_t = u_t + \beta\delta u_{t+1} + \beta\delta^2 u_{t+2} + \beta\delta^3 u_{t+3} + \cdots$$

$$= u_t + \beta(\delta u_{t+1} + \delta^2 u_{t+2} + \delta^3 u_{t+3} + \cdots)$$

Quasi-hyperbolic discounters choose more patiently in the long than in short run:

$$\text{Future discount rate} = \frac{D(\tau) - D(\tau+1)}{D(\tau)} = \frac{\beta\delta^{\tau} - \beta\delta^{\tau+1}}{\beta\delta^{\tau}} = 1 - \delta \approx 0.05$$

$$\text{Present discount rate} = \frac{D(0) - D(1)}{D(0)} = \frac{1 - \beta\delta}{1} = 1 - \beta\delta \approx 0.5$$

Time-inconsistency

When $\beta < 1$, the quasi-hyperbolic β, δ model can generate a conflict between earlier and later selves and the time-inconsistency it causes.

The time-inconsistent model predicts "self-control problems" like procrastination.

For the examples that follow, consider a special case to build intuition: $\beta = 1/2$ and $\delta = 1$:

$$D(\tau) = \{1, \beta\delta, \beta\delta^2, \beta\delta^3, \ldots\} = \{1, ½, ½, ½, \ldots\}.$$

- Relative to the present period, all future periods are worth less.
- All discounting takes place between present and immediate future.
- In the 'long run' we're relatively patient–utils tomorrow are just as valuable as utils the day after tomorrow.

- Exercise has benefit today of -6. Exercise has delayed benefit of 8.

- Exercise today? No.

$$-6 + \frac{1}{2}8 = -2 < 0$$

- Exercise tomorrow? Yes.

$$0 + \frac{1}{2}(-6 + 8) = 1 > 0$$

- But tomorrow you'll again want to postpone exercising (Akerlof 1991 *AER*; O'Donoghue and Rabin 1999 *AER*; Dellavigna and Malmendier 2006 *AER*)

The quasi-hyperbolic model implies that decisions are sensitive to present vs. future timing:

Are the utility benefits and costs simultaneous, or do benefits precede costs?

• Planning to eat candy tomorrow (simultaneous benefits and costs)

$$\beta\delta(B_{Pleasure} - \delta C_{Health}) = \frac{1}{2}(2 - 3) < 0$$

• Eating candy today (benefits precede costs)

$$(B_{Pleasure} + \beta\delta(-C_{Health}) = 2 + \frac{1}{2}(-3) > 0$$

Wertenbroch (1998 *Marketing Science*) found that consumers tended to buy "temptation goods" in small packages, foregoing volume discounts: exhibiting a demand for commitment that they would not need without present bias.

• Buying future *People* magazines by subscribing (simultaneous benefits and costs)

$$\beta\delta\left(B_{Pleasure} - C_{Money}\right) = \frac{1}{2}(2-3) < 0$$

• Buying present *People* magazine from check-out line (benefits precede costs)

$$\left(B_{Pleasure} + \beta\delta(-C_{Money}\right) = 2 + \frac{1}{2}(-3) > 0$$

(Note that the financial cost is later, since as a natural consequence of the dynamic budget constraint current expenditure crowds out later expenditure.)

Oster and Scott-Morton (2005 *BE Advances in Economic Analysis & Policy*) found that:

• Magazines like *People* sold on newsstands at high price relative to subscription

• Magazines like *Foreign Affairs* sold on newsstands at low price relative to subscription

• But, disproportionately, magazines like *People* are sold on newsstands and magazines like *Foreign Affairs* are sold by subscription.

- Order future highbrow movie instead of future fun movie (simultaneous benefits and costs)

$$\beta\delta\left(B_{Knowledge} - B_{Fun}\right) = \frac{1}{2}(3-2) > 0$$

- Watch present highbrow movie instead of present fun movie (benefits lag costs)

$$-B_{Fun} + \beta\delta B_{Knowledge} = -2 + \frac{1}{2}(3) < 0$$

The details matter (much more than for time-consistent choice):

- Build up $5000 of debt on a credit card at 20% interest? Yes.

- Take out a home equity loan at 5% interest requiring three hours of paperwork and a two week processing delay? I'll do it next week.

- Take out a home equity loan at 10% interest—'Preapproved with No Paperwork Required'? Yes.

  (But what if the potential borrower who receives the 10% offer knows that 5% loans are available at another bank?)

- Buy a new car, making $2000 down-payment? No.

- Buy a new car, paying more interest, but making no down payment (e.g. "$2000 cash back")? Yes.

Inferences from static choices: illustrative model

Recall Read and van Leeuwen's (1998 *Organizational Behavior and Human Decision Processes*) inferences from static choices in kind (e.g. chocolate at t or orange at t):

- Overwhelming majority of subjects chose a healthy snack now to eat next week.
- Overwhelming majority of subjects chose an unhealthy snack now to eat now.

Let c represent calories (cocaine, candy, cigarettes), with u" concave and u" < v".

$$U(c_{t-1}, c_t, c_{t+1}, \dots) = [u(c_t) - v(c_{t-1})] + \beta\delta[u(c_{t+1}) - v(c_t)] + \beta\delta^2[u(c_{t+2}) - v(c_{t+1})],$$

Optimal level of consumption from the perspective of date t:

$$u'(c_t) = \beta\delta v'(c_t)$$

$$u'(c_{t+1}) = \delta v'(c_{t+1})$$

$$\beta = 1 \text{ if and only if } c_t = c_{t+1}$$

$$\beta < 1 \text{ if and only if } c_t > c_{t+1}$$

(If $\beta = 1$, all c_t satisfy the same first-order condition. β does not influence c_{t+1}. Hence $\frac{dc_t}{d\beta} = \frac{\delta v\prime}{u"-\beta\delta v"} < 0$, given that the denominator is negative by the second-order conditions. Thus $c_t > c_{t+1}$ implies $\beta < 1$.)

Empirical application

Laibson (1997 *QJE*) applies present-biased preferences to savings-consumption decisions.

In particular, he explores the role of illiquid assets as an imperfect commitment technology in saving-consumption decisions.

Illiquid assets: Assets for which either (i) there is a period of time before which you can convert them to cash, or (ii) there is a penalty if you convert them to cash.

For instance:

- Long term: retirement plans
- Medium term: house, durable goods, business
- Short term: CD accounts.

Model

Consumer makes decisions in periods 1, ...,T.

One liquid asset x and one illiquid asset z.

Exogenous initial asset holdings x_0, $z_0 \geq 0$.

In period t:

- Earn labor income y_t
- Earn asset income $R_t(x_{t-1} + z_{t-1})$
- Choose consumption c_t and new asset allocation x_t and z_t such that

$$c_t + x_t + z_t = y_t + R_t(x_{t-1} + z_{t-1})$$

$$c_t \leq y_t + R_t x_{t-1}$$

$$x_t, z_t \geq 0$$

Period-t intertemporal preferences:

$$U^t = E_t [u(c_t) + \beta \sum_{\tau=t+1}^{T} \delta^{\tau-t} u(c_\tau)]$$

Assumes sophistication.

Example:

- Suppose $x_0 = W$, $z_0 = 0$, and $y_t = 0$ for all t.
- Assume $\delta = 1$ and $R_t = 1$ for all t.

For $\beta = 1$ (Time Consistents):

- Optimum involves $c_t = W/T$ for all t.
- To implement, choose $x_t = W - t(W/T)$ for all t.

Example:

- Suppose $x_0 = W$, $z_0 = 0$, and $y_t = 0$ for all t.
- Assume $\delta = 1$ and $R_t = 1$ for all t.

For Sophisticates with $\beta < 1$ (quasi-hyperbolic):

- Desired behaviour from period-1 perspective:

  $c_1 = \overline{c}$ and $c_t = \underline{c}$ for all $t > 1$, where $u'(\overline{c}) = \beta u'(\underline{c})$ and $\overline{c} + (T-1)\underline{c} = W$.

- Can you implement it? Yes.

  $c_1 = \overline{c}$, $x_1 = \underline{c}$, $z_1 = W - \overline{c} - \underline{c}$

  $c_2 = \underline{c}$, $x_2 = \underline{c}$, $z_2 = W - \overline{c} - 2\underline{c}$

  $c_t = \underline{c}$, $x_t = \underline{c}$, $z_t = W - \overline{c} - t\underline{c}$.

Example: T = 3

- Suppose $x_0 = z_0 = 0$ and $y_1 > y_2 > 0 = y_3$ (define $y_1 + y_2 = W$).

- Assume $\delta = 1$ and $R_t = 1$ for all t.

For $\beta = 1$ (Time Consistents):

- Optimum involves $c_t = W/3$ for all t.

- Implement as before.

Example: T = 3

• Suppose $x_0 = z_0 = 0$ and $y_1 > y_2 > 0 = y_3$ (define $y_1 + y_2 = W$).

• Assume $\delta = 1$ and $R_t = 1$ for all t.

For Sophisticates with $\beta < 1$:

• "Desired" behavior from period-1 perspective:

$c_1 = \overline{c}$ and $c_2 = c_3 = \underline{c}$, where $u'(\overline{c}) = \beta u'(\underline{c})$ and $\overline{c} + 2\underline{c} = W$.

• Can you implement it?

If $y_2 \leq \underline{c}$, we can implement it much as before:

$c_1 = \overline{c}$, $x_1 = \underline{c} - y_2$, $z_1 = \underline{c}$.

Example: T = 3

- Suppose $x_0 = z_0 = 0$ and $y_1 > y_2 > 0 = y_3$ (define $y_1 + y_2 = W$).
- Assume $\delta = 1$ and $R_t = 1$ for all t.

For Sophisticates with $\beta < 1$:

- "Desired" behavior from period-1 perspective:

  $c_1 = \overline{c}$ and $c_2 = c_3 = \underline{c}$, where $u'(\overline{c}) = \beta u'(\underline{c})$ and $\overline{c} + 2\underline{c} = W$.

- Can you implement it?

If $y_2 > \underline{c}$, we cannot implement it because period-2 self will consume more than $\underline{c}$.

But still might use the illiquid asset:

- Fixing W, there exists $\overline{y} > \underline{c}$ such that for all $y_2 \in (\underline{c},\overline{y})$, the outcome is $c_1 = \overline{c}(y_2)$, $c_2 = y_2$, and $c_3 = \underline{c}(y_2)$, where $u'(\overline{c}(y_2)) = \beta u'(\underline{c}(y_2))$ and $\overline{c}(y_2) + \underline{c}(y_2) = W - y_2$.
- Implement by choosing $x_1 = 0$ and $z_1 = \underline{c}(y_2)$.

Implications

- Hyperbolic discounters endogenously tie up their wealth in illiquid assets (avoid holding wealth in liquid assets).

- Hyperbolic discounters will exhibit comovement of consumption and income.

- For hyperbolic discounters, financial innovations that reduce illiquidity (e.g., credit cards) may decrease the savings rate and decrease “welfare”.

Another empirical application

Laibson, Repetto, and Tobacman (2008 *Brookings Papers on Economic Activity*) (see also Angeletos et al. 2001 *JEP*) used method of simulated moments to estimate discounting parameters in a life-cycle model of income, consumption, and credit-card borrowing.

They used four moments:

- Proportion of households carrying a credit-card balance.
- Average over the life cycle of

  mean credit-card borrowing for age x
  mean income for age x

- MPC out of expected income changes.
- Wealth-to-income ratios for households aged 50-59.

Findings

Substantial illiquid retirement wealth: W/Y = 3.9.

Extensive credit card borrowing:

- 68% didn't pay their credit card in full last month
- Average credit card interest rate is 14%
- Credit card debt averages 13% of annual income

Consumption-income comovement:

- Marginal Propensity to Consume = 0.23 (i.e. consumption tracks income)

Simulation Model

- Stochastic Income
- Lifecycle variation in labor supply (e.g. retirement)
- Social Security system
- Life-cycle variation in household dependents
- Bequests
- Illiquid asset
- Liquid asset
- Credit card debt

Numerical solution (via backward induction) of 90-period lifecycle problem.

Results

$$U_t = u_t + \beta[\delta u_{t+1} + \delta^2 u_{t+2} + \delta^3 u_{t+3} + ...]$$

- $\beta = 0.70$ (standard error 0.11)
- $\delta = 0.96$ (standard error 0.01)
- Null hypothesis of $\beta = 1$ rejected (t-statistic of 3).
- Specification test accepted.

Moments:

	Empirical	Simulated (Hyperbolic)
%Visa:	68%	63%
Visa/Y:	13%	17%
MPC:	23%	31%
f(W/Y):	2.6	2.7

Intuition

- Long run discount rate is $-\ln(\delta) = 4\%$, so save in long-run (illiquid) assets.

  (Recall: To relate a discount factor δ to a discount rate ρ, $\rho = -\ln \delta$.)

- Short-run discount rate is $-\ln(\beta\delta) = 40\%$, so borrow on your credit card today.

- Indeed, you might even borrow on your credit card so you can "afford" to save in your 401(k) account.

Another empirical application

Shapiro (2005 *J. Public Economics*) examined consumption patterns among food-stamp recipients, and found that over the month between food-stamp deliveries, caloric intake declines by about 10-15 percent.

- Survey evidence revealed rising desperation over the month, suggesting that high elasticity of intertemporal substitution is not a likely explanation.

- Households with more short-run impatience (estimated from hypothetical intertemporal choices) were more likely to run out of food sometime during the month.

- To explain this behavior with a standard exponential-discounting model would require calibration with an extreme annual discount factor of 0.23.

- By contrast, a quasi-hyperbolic discounting model gracefully explains the data with a plausible daily $\beta = 0.91$ and $\delta = 1$.

The data can reject a number of alternative hypotheses:

- Households that shop for food more frequently do not display a smaller decline in intake over the month, casting doubt on depreciation stories.

- Individuals in single-person households experience no less of a decline in caloric intake over the month than individuals in multi-person households.

- Survey respondents are not more likely to eat in another person's home toward the end of the month.

- The data show no evidence of learning over time.

Naifs and Sophisticates

Some of the phenomena discussed above, e.g. Ariely and Wertenbroch's and Dellavigna and Malmendier's, cannot be explained by present bias alone.

People have therefore considered models that combine present bias with less-than-perfect sophistication about the person's own future behavior.

(Odysseus was as remarkable for his sophistication about his own behaviour (e.g. the sirens) as he was for his sophistication about others' (the horse).)

• Naifs falsely believe future selves will maximize today's preferences (Strotz 1957 *REStud*).

—Solution concept: maximization (mispredict future discount rates)

—Prediction (Dellavigna and Malmendier 2006 *AER*): never exercise (but join gym).

• Sophisticates have rational expectations (Strotz 1957 *REStud*).

—Solution concept: subgame perfect equilibrium.

—Prediction: never exercise (and don't join gym).

• Partial naivete (O'Donoghue and Rabin 2001 *QJE*)

—Solution concept: subgame perfect equilibrium, using $\hat{\beta}$ such that $\beta < \hat{\beta} < 1$.

(Note that naifs use $\hat{\beta} = 1$ and sophisticates use $\hat{\beta} = \beta$.)

Illustrative model (O'Donoghue and Rabin 1999 *AER*)

- Must do a project in one of T < ∞ periods, at a cost that increases over time.

- In time period t the project costs $(\frac{3}{2})^t$ utils to execute. (No discounting.)

- Each period, must choose to either "do" or "wait"; commitment is impossible.

- If wait until period T, must do it then.

With no discounting, an exponential discounter would do it now, with or without commitment.

Naifs

Naifs are fully unaware of their future self-control problems and therefore expect to behave in future exactly as they currently would like to behave in the future.

• When will a quasi-hyperbolic naïf with β = 1/2 and δ = 1 so the project?

From the current self's perspective, it's always better to postpone doing the project until next period:

$$\left(\frac{3}{2}\right)^t > \beta\delta\left(\frac{3}{2}\right)^{t+1} = \frac{1}{2}\left(\frac{3}{2}\right)^{t+1} = \frac{3}{4}\left(\frac{3}{2}\right)^t < \left(\frac{3}{2}\right)^t$$

(Even partial naifs can make the same kind of mistake.)

Sophisticates

Sophisticates are quasi-hyperbolic with $\beta < 1$ hence time-inconsistent, but are aware of their self-control problems and so correctly predict their future behavior.

Hence they follow a finite-horizon subgame-perfect equilibrium, treating their own future selves as independent players.

- When will a quasi-hyperbolic sophisticate with $\beta = 1/2$ and $\delta = 1$ do the project?

If T is even, sophisticates will do the project in even but not in odd periods (complete contingent plan, as required by the strategic analysis):

- In the penultimate period, it's better to postpone for the same algebraic reason that it was better for a naïf, because a sophisticate believes (now correctly) that his period-T self will do the project. And so on, working back to the start.

If T is odd, then sophisticates will do the project in odd but not in even periods, for similar reasons.

The model's implications are less bizarre with uncertainty or an infinite horizon.

Another problem with the sophisticated (or partially naive) model is that the model predicts that people will often seek out and pay for commitment.

But:

- We see surprisingly little endogenous commitment for commitment's sake.

- Most commitment is ancillary (e.g. obligatory monthly mortgage payments).

- Very little is gratuitous and advertised as such (e.g. Christmas clubs).

One interesting exception: the new commitment website StickK.

Due to behavioural economists Dean Karlan and Ian Ayres, so may not “count”.

Another, Clocky, an alarm clock that plays annoying tunes while running away from you, unless you catch it and turn it off http://www.nandahome.com/:

Still others (“Get up now! Or a donation to [your favourite odious charity] will be automatically deducted from your account”):

Experiments on commitment

In a field experiment, Ashraf, Karlan, and Yin (2006 *QJE*):

- Offered a commitment savings product to randomly chosen clients of a Philippine bank

- 28.4% take-up rate of commitment product (either date-based goal or amount-based goal)

- More “hyperbolic” subjects were more likely to take up the product

- After twelve months, average savings balances increased by 81% for those clients assigned to the treatment group relative to those assigned to the control group.

In a field experiment, Gine, Karlan, and Zinman (2010 *AEJ Applied*):

- Tested a voluntary commitment product (CARES) for smoking cessation.
- Smokers offered a savings account in which they deposit funds for six months, after which take urine tests for nicotine and cotinine.
- If they pass, money is returned; otherwise, forfeited
- 11% of smokers offered CARES take it up, and smokers randomly offered CARES were 3 percentage points more likely to pass the 6-month test than the control group
- Effect persisted in surprise tests at 12 months.

Applied theory

Much experience with credit cards and subprime mortgages suggests that many consumers overborrow, do not comply fully with repayment terms and pay large penalties for small deviations from compliance.

This would seldom happen in a neoclassical competitive credit market.

Present-bias seems to be part of the story, because the initial terms offered borrowers are usually much more favorable than the eventual terms.

But present-bias alone is not enough, because much such borrowing, e.g. to purchase a durable good, has up-front effort costs and mostly delayed benefits.

Heidhues and Koszegi (2010 *AER*) analyzed contract choices, loan-repayment behavior, and consumer welfare in a model of a competitive credit market when borrowers are both present-biased and naive.

Their model has three periods, 0, 1, and 2.

If the consumer borrows $c \geq 0$ in period 0 and repays $q \geq 0$ and $r \geq 0$ in periods 1 and 2, then self 0 has utility $c - k(q) - k(r)$, where $k(\bullet)$ is the cost of repayment.

Self 1 maximizes $-k(q) - \beta k(r)$ for $0 < \beta \leq 1$, so for $\beta < 1$ he is present-biased.

(Self 2 makes no decisions. His only role in the model is to repay, and suffer….)

Note that self 0 is *not* present-biased in favor of c over reducing k(r).

This captures the idea that most borrowing is for future consumption:

When you take out a mortgage to buy a house, you do get to move in in a few weeks or months, and this might undermine your self-control. But the evidence suggests that to create a self-control problem, immediate gratification must be more immediate; and in any case, most house benefits will be further delayed.

The consumer may be sophisticated, naïve, or partially naïve as in O'Donoghue and Rabin 2001 *QJE*.

Thus self 0 believes with certainty that self 1 will maximize - k(q) - $\hat{\beta}$k(r) for some $\beta \leq \hat{\beta} \leq 1$.

$\hat{\beta} = \beta$ represents perfect sophistication and $\hat{\beta} = 1$ represents complete naivete regarding time inconsistency.

(Heidhues and Koszegi allow heterogeneous naivete, but that's less important.)

The market is perfectly competitive, with firms knowing everything about consumers and consumers knowing everything about firms and themselves, except that less than fully sophisticated consumers don't know their true βs.

(This is a natural modeling choice, because many people think that perfect competition is a panacea. But the results would be similar for a monopoly.)

In this context a competitive market is cleared by contracts in period 0, with consumers' 0 selves choosing among the contracts offered by firms; or choosing no contract, which yields them an exogenous reservation utility.

In equilibrium (specializing Heidhues and Koszegi's Definition 2 to homogeneous consumers), each consumer chooses the offered contract that seems optimal in period 0, each firm makes 0 expected profit, no firm can earn positive profit by deviating to another contract, and in period 1, without regard to their sophistication, consumers behave according to their actual βs.

(Thus we solve via backward induction in any case; but only for sophisticated consumers is behavior in subgame-perfect equilibrium among the selves, and only for sophisticated consumers is behavior time-consistent.)

For simplicity Heidhues and Koszegi restrict attention to non-redundant contracts. We can then think of a firm as selecting consumption c along with a “baseline" repayment schedule that self 0 expects to choose (and a sophisticated consumer would actually choose), plus the alternative repayment schedule that a naïve consumer actually chooses.

In equilibrium each firm chooses the contract that maximizes its profit, given how the consumer actually behaves, subject to the constraints that:

- self 0 weakly prefers the baseline repayment schedule (the one that self 0 expects to choose) to his reservation utility, as in a standard participation constraint, but with incorrect expectations; and

- self 0 expects that he will choose the baseline repayment schedule, and must therefore expects self 1 to prefer it to the alternative schedule, as in a standard incentive-compatibility constraint.

These contracts must also yield zero profit. Otherwise a firm could do better by slightly improving its contract terms and competing away all consumers.

Heidhues and Koszegi obtain several results.

- Both time-inconsistency and naivete are needed for the explanation to work: With time-consistent and sophisticated consumers, or even with time-consistent and naïve consumers, equilibrium contracts maximize welfare, and rules against abusive credit practices are nonbinding.

- For time-inconsistent and less than perfectly sophisticated borrowers, the profit-maximizing contract has cheap baseline repayment terms, but the terms are inefficiently front-loaded and noncompliance incurs a large penalty. The terms are inefficient because the firm chooses them to appeal to self 0, whose naivete distorts the trade-offs. The penalties are also inefficient, but a naïve consumer doesn't expect to pay them, and this distorts the trade-offs.

- Naïve consumers borrow more than sophisticated consumers.

- Although credit is for future consumption, naïve consumers end up back-loading repayment, incurring large and unanticipated costs.

As in most of the literature, Heidhues and Koszegi take the long-term perspective on welfare and equate the consumer's welfare with self 0's utility.

- By this standard naïve consumers have lower welfare than sophisticated ones: discontinuously lower, no matter how close to sophisticated they are.

Many people think that competition will police contractual markets.

This is true in the model with time-consistent consumers, even if they are naïve.

But it fails badly with even slightly naïve time-inconsistent consumers. Competitive firms *must* exploit such consumers or be priced out of the market.

Heidhues and Koszegi show that prohibiting large penalties for deferring small amounts of repayment—akin to recent regulations—can raise welfare.

Importantly, such regulations can benefit naïve consumers without affecting sophisticated consumers: what some analysts call “libertarian paternalism”.

More examples of applied theory

Naive procrastination (various papers by O'Donoghue and Rabin).

Sophisticated information acquisition — sophisticates' incentives to acquire information will be distorted if that information is likely to alter future misbehavior (Carrillo and Mariotti, 2000 *REStud*, Benabou & Tirole, 2002 *QJE*).

Consumption of addictive products (Gruber and Koszegi, 2001 *QJE*, O'Donoghue and Rabin, 1999 Book Chapter).

Optimal taxes for "sin" goods (O'Donoghue and Rabin, 2006 *J. Public Economics*).

Job search (DellaVigna & Paserman, 2005 *J. Labor Economics*).

Smoking and implications for the effects of cigarette taxes (Gruber and Mullainathan 2005 *BE Advances in Economic Analysis & Policy*).

General lessons on procrastination

- Severe procrastination if and only if the person is at least partially naive.
  (A sophisticated person knows enough to limit the costs.)

- But there is no corresponding limit for a naïve person who never learns, and the welfare implications of even mild self-control problems and mild naivete can be large.

- Higher stakes may not improve behavior: In some contexts, the more important something is, the more likely we are to procrastinate.

- One cannot infer long-run preferences from choice behavior.
  (Without considering short-run incentives, there is no basis for concluding that the connection is close.)

- Details matter.