Lecture 2: Dynamically Inconsistent Preferences

Ran Spiegler

August 26, 2007

Abstract

Several psychological factors cause decision makers' tastes to change over time. In this lecture we introduce the multi-selves approach to modeling decision making when tastes are changing, with various assumptions on the decision maker's ability to anticipate the change in his tastes. We demonstrate this modeling tool with a market application, in which a profit-maximizing firm interacts with a dynamically inconsistent consumer.

Choice behavior - captured by a choice function over a set of consequences X - is rational if it is consistent with maximization of a preference relation $\succsim$ over X. If the decision maker's choices sometimes reveal that he prefers x to y and sometimes reveal that he prefers y to x, then he lacks stable preferences and in this sense, his choice behavior is not rational.

In many economic situations of interest, it makes sense to assume that such violations of consistency have a temporal regularity. That is, in the course of a dynamic decision problem, the decision maker's preferences change in a systematic way. Here are a few examples.

Present bias. Recall the following experiment from the previous lecture. A consequence is a pair consisting of an amount of money m and the calendar day t in which the amount is obtained. In the experiment, people are faced with the problem of choosing between such pairs (m, t). Let $t > 0$. On day 0, experimental subjects tend to prefer $(11, t+1)$ to $(10, t)$. On day t itself, however, they tend to prefer $(10, t)$ to $(11, t+1)$. That is, people are willing to wait an extra day for an increased amount of money when the smaller amount is scheduled to arrive in the future, yet when the smaller amount is available immediately, they prefer not to wait.

Temptations. Consider a two-stage decision problem. At the first stage, you choose a restaurant. At the second stage, while at the restaurant, you decide which dish to order. The set of consequences X is simply the set of dishes. You are on a diet, so from the point of view of the first stage, eating steak is inferior to eating salad. If you could, you would commit to eat salad. However, once at the restaurant, you are tempted by the steak and go for it.

Dynamic implications of reference-point effects. Suppose that when you do not own an object, you are willing to pay at most \$20 for it. But when you own the object, you are willing to sell it for no less than \$25. One reason for such a discrepancy between willingness to pay and willingness to accept is that people tend to attach a special significance to the status quo, and they require special reasons for changing it. A consequence in this example is a specification of whether you own the object, as well as the amount of money at your disposal. Imagine that at date 0, you do not possess the object. You are faced with the following business opportunity: purchasing the object for \$22 and then re-sell it at date 1 for \$24. From the point of view of date 0, you prefer the final outcome of having a net monetary gain of \$2 without having the object to the final outcome of owning the object with a net monetary loss of \$22. From the point of view of date 1, *conditional on having purchased the object*, your preference is reversed.

Note that in the first example, the change in preferences is due to the mere passage of time. In the other two examples, preferences change as a result of a particular contingency that may arise in the course of a dynamic decision problem. In the second example, it is the experience of actually sitting in a restaurant in front of a menu (and perhaps seeing other people having steak). In the third example, it is the state of owning the object.

These examples suggest a large variety of economic situations in which changing tastes could be relevant: life-cycle saving and consumption, trading in exchange markets, consumption of addictive goods, etc. Thus, it would be nice to have a modeling tool that enables us to say some things about choice behavior under changing tastes. On one hand, the inconsistent preferences over X displayed by decision makers in the above examples imply that we are unable to account for their behavior with the model of maximization according to a preference relation over X. However, the fact that the dynamic inconsistencies are systematic means that we may be able to find alternative modeling tools to account for them. In this lecture, we introduce such a modeling tool and apply it to an economic situation of interest.

1 The Multi-Selves Approach

The most widely practiced modeling tool that economists use to analyze decision making under changing tastes is known as the "multi-selves approach". Given an extensive-form decision problems, we model the decision maker as a collection of different players ("selves") having idiosyncratic preferences. We assign a self to a decision node, in the same way that we would assign a player to a node in an extensive-form game. Thus, instead of modeling the decision maker's choices as the outcome of maximization, we model it as the outcome of some game-theoretic solution concept, in a game played among the different selves. Note that this modeling approach involves making assumptions on how each self forms beliefs regarding the future actions of other selves.

The simplest application of this methodology is to two-period decision problems of the type captured by the restaurant example. In period 1, the decision maker chooses a menu, namely a non-empty subset of X. In period 2, he chooses an element from the chosen menu. This is modeled as a Stackelberg game, such that self j moves in period j, $j = 1, 2$, and self j's preferences over X are $\succsim_j$. The "standard" solution concept for Stackelberg games is subgame perfect *equilibrium*. Applied in this context, this means that the decision maker can perfectly anticipate the future change in his tastes.

This model clearly explains the taste for commitment that people often display in situations such as the restaurant example. For instance, let $X = \{steak, lettuce\}$, $lettuce \succ_1 steak$ and $steak \succ_2 lettuce$. Then, in subgame perfect equilibrium, self 1 will choose the menu $\{lettuce\}$ in period 1, and consequently, self 2 will be forced to eat lettuce in period 2. In equilibrium, self 1 strictly prefers $\{lettuce\}$ to any other menu - in particular, to the grand set X. This strict taste for commitment is a general real-life phenomenon. For instance, knowing that you will be tempted to eat an entire packet of cashew nuts, you choose to constrain your future options by buying a small packet. Similarly, knowing that you will be tempted to squander your entire disposable income on fun and games, you lock a substantial portion of your income in an illiquid savings account, in order to deter yourself from expending too much on fun and games.

The possibility of a strict preference for commitment is a behavioral implication of the multi-selves model which clearly distinguishes it from a standard decision model, applied to the same class of two-period decision problems, in which the decision maker has stable preferences over X. In this sense, the multi-selves model is a highly successful model.

What are the behavioral implications of the multi-selves model for first-period behavior, apart from the general taste for commitment? It turns out that we are able to fully characterize it with a simple axiom. The multi-selves model implies a preference relation that the decision maker has over menus in the first period. Formally, assume that X is finite, let $P(X)$ be the set of all choice problems - i.e., non-empty subsets of X, and let $\succsim^*$ be a preference relation over $P(X)$. The following is an axiom that $\succsim^*$ may satisfy:

Axiom 1 *For every $A, B \in P(X)$, $A \sim^* A \cup B$ or $B \sim^* A \cup B$.*

Note that this axiom accommodates the taste for commitment described above, because it allows $A \cup B$ to be equivalent to the inferior menu among A and B. However, it is clear that the axiom imposes a lot more structure on the taste for commitment.

Proposition 1 *A preference relation $\succsim^*$ over $P(X)$ satisfies Axiom 1 if and only if there exist two numerical functions, $u : X \to \mathbb{R}$ and $v : X \to \mathbb{R}$, such that for every $A, B \in P(X)$,*

$$A \succ^* B \quad \textit{if} \max_{x \in \arg\max_{x \in A} v(x)} u(x) > \max_{y \in \arg\max_{x \in B} v(x)} u(y)] \tag{1}$$

Thus, Axiom 1 completely characterizes the first-period behavior induced by the two-period multi-selves model. Instead of proving this result, let us prove a slightly easier one. Note that the axiom immediately implies that for every menu A, there exists some $x \in A$ such that $A \sim^* \{x\}$. The following axiom strengthens this property.

Axiom 2 *For every $A \in P(X)$, there is a* ***unique*** *element $x \in A$, such that $A \sim^* \{x\}$.*

The following is a specification of the multi-selves model, in which the two numerical functions u and v are one-to-one.

Proposition 2 *A preference relation $\succsim^*$ over $P(X)$ satisfies Axioms 1 and 2 if and only if there exist two* ***one-to-one*** *numerical functions, $u : X \to \mathbb{R}$ and $v : X \to \mathbb{R}$, such that for every $A, B \in P(X)$,*

$$A \succ^* B \quad \textit{if } u[\arg\max_{x \in A} v(x)] > u[\arg\max_{x \in B} v(x)] \tag{2}$$

Proof. First, let us assume that $A \succ^* B$ if and only if $u[\arg\max_{x\in A} v(x)] > u[\arg\max_{x\in B} v(x)]$, and show that $\succsim^*$ satisfies Axioms 1 and 2. To see that Axiom 1 is satisfied, note that $\arg\max_{x\in A\cup B} v(x)$ belongs to A or to B. Therefore, $\arg\max_{x\in A\cup B} v(x) = \arg\max_{x\in A} v(x)$ or $\arg\max_{x\in A\cup B} v(x) = \arg\max_{x\in B} v(x)$, hence $A \cup B \sim^* A$ or $A \cup B \sim^* B$. Axiom 2 follows immediately from the assumption that both u and v are one-to-one functions.

Second, assume that $\succ^*$ satisfies Axioms 1 and 2, and let us construct the one-to-one functions u and v, such that $A \succ^* B$ if and only if $u[\arg\max_{x\in A} v(x)] > u[\arg\max_{x\in B} v(x)]$. Construct u as follows: $u(x) \geq u(y)$ if $\{x\} \succsim^* \{y\}$. Because $\succsim^*$ is a preference relation, u is well-defined. Let us show that u is a one-to-one function. Assume the contrary. Then, there exist $x, y \in X$ such that $\{x\} \sim^* \{y\}$. By Axiom 1, $\{x, y\} \sim^* \{x\}$ or $\{x, y\} \sim^* \{y\}$. Therefore, $\{x, y\} \sim^* \{x\} \sim^* \{y\}$, contradicting Axiom 2.

Let us turn to constructing v. To do so, define a choice function c. For every menu $A \subseteq X$, $c(A)$ is the element $x \in A$ for which $\{x\} \sim^* A$. By Axiom 2, there is exactly one such element. Therefore, c is a well-defined choice function. Let us show that c satisfies IIA. Let $c(A) = x$, and consider a subset $B \subset A$ which contains x. Let us show that $\{x\} \sim^* B$. Assume the contrary. Then, $A \nsim^* B$. By Axiom 1, $A \sim^* A \backslash B$. By Axiom 2, there exists an element $y \in A \backslash B$ such that $\{y\} \sim^* A \backslash B$. Therefore, $\{y\} \sim^* A$. But since by assumption $c(A) = x \neq y$, and since by definition this means that $\{x\} \sim^* A$, we get a violation of Axiom 2. Therefore, $\{x\} \sim^* B$, hence $c(B) = x$. We have established that c satisfies IIA. Therefore, c can be rationalized by a one-to-one utility function v over X.

Suppose that $A \succ^* B$. By the definition of v, $A \sim^* \{\arg\max_{x\in A} v(x)\}$ and $B \sim^* \{\arg\max_{x\in B} v(x)\}$. By the definition of u, $\{\arg\max_{x\in A} v(x)\} \succ^* \{\arg\max_{x\in B} v(x)\}$ if and only if $u[\arg\max_{x\in A} v(x)] > u[\arg\max_{x\in B} v(x)]$. This concludes the proof. ■

The axiomatic characterization of first-period behavior implied by the multi-selves model deepens our understanding of the model, and sometimes helps us analyze applications without getting into laborious calculations.

2 Criticisms of the Multi-Selves Approach

When we turn from the above class of two-period decision problems to more general classes of decision problems, the multi-selves approach becomes problematic for several reasons. First, applying subgame perfection to extensive games with a long horizon is known to generate paradoxical results, because most people do not apply the logic of backward induction for more than three iterations.

Second, there are ambiguities in the approach. Consider a completely standard, single-self dynamic decision model. On each day, you decide whether to clean your office in the morning and in the evening. Your payoff on day t is 3 if you clean your office twice, 0 if you do it once, and 1 if you do not do it at all. Your overall utility is a discounted sum of periodic payoffs. Your long-run utility is a discounted sum of periodic payoffs. Now, instead of modeling the situation as a single-self decision problem, let us model it as a repeated game between your "morning self" and your "evening self". Both selves have the same payoff function defined over terminal histories of the game as in the single-self version. Thus, the situation becomes an infinitely repeated 2×2 coordination game. There is no dynamic inconsistency, since both selves have the same preferences. In particular, they both agree that the best thing to do is clean the office twice a day indefinitely. However, the game has Pareto inefficient subgame perfect equilibria - for instance, no self ever cleans the office. We could use the Pareto selection criterion to rule out inefficient equilibria. However, we do not normally use this criterion in multi-player games. Why should we feel more comfortable doing it in an intrapersonal game?

A third criticism of the multi-selves approach has to do with the monitoring technology. In standard game-theoretic analysis, we assume that at most, a player can monitor the past actions of his opponents (as well as his own actions). He certainly cannot monitor their *motivations*. In particular, he cannot tell whether an action was taken deliberately or by mistake. Game-theoretic solution concepts are constructed on the basis of this inability to distinguish between deliberate and accidental actions. However, a single decision maker does have access to his own intentions, and can therefore tell whether a past action was deliberate or accidental. Given this fundamental difference, should we devise new solution concepts that are specifically sensitive to it?

2.1 Naivety

As mentioned earlier, subgame perfect equilibrium implies that the decision maker is fully aware of the future change in his preferences. In other words, he is "sophisticated". In many situations, it seems reasonable to assume that people underestimate the likelihood that (or the extent to which) their tastes may change over time or as a result of a change in circumstances. For example, people may underestimate the extent to which changing their position in an organization will affect their trade-off between work and leisure; they are often overconfident about their ability to resist temptation and maintain a diet; etc.

In the two-period model described above, we can capture naivety by assuming that in period 1, self 1 has a prior belief over self 2's preferences, which assigns probability α to $\succsim_1$ and probability $1-\alpha$ to $\succsim_2$. Thus, α is the decision maker's subjective probability that his tastes will not change. More generally, we may assume that self 1 has a subjective probability distribution over second-period preferences, whose support consists of preference relations which are in some sense "between" $\succsim_1$ and $\succsim_2$.

In the two-period model, suppose that self 1 is "fully naive" - i.e., he is certain that his preferences over X will not change. Then, his choice of menus in period 1 will be indistinguishable from that of a standard decision maker with stable preferences over X. However, his second-period choice will belie such an interpretation: although in period 1 his induced preferences over menus will be, say, $\{steak, lettuce\}\tilde{}^*\{lettuce\} \succ^* \{steak\}$, his second-period choice from the menu $\{steak, lettuce\}$ will be $steak$.

2.2 An Example: Procrastination

There is a task that a decision maker needs to carry out once and for all. He is entitled to choose to perform it in any period $k \in \{1, ..., T\}$. Every period is assigned a different self. The preferences of self t over the subset $\{t, ..., T\}$ are $t+1 \succ_t t \succ_t t+2 \succ_t t+3 \succ_t \cdots \succ_t T$. That is, self t prefers to postpone performing the task to the next period, but not more than that. The psychology behind such preferences is that the decision maker has a present bias. As far as future periods are concerned, the decision maker knows that it is better for him to complete the task earlier (due to the mounting material cost of not doing it). However, when the choice is between doing it today and doing it tomorrow, he prefers to procrastinate.

Assume that the decision maker is fully aware of his future tastes. What is the subgame perfect equilibrium of the game played by the multiple selves? Backward induction provides a simple answer. At period $T-1$, the current self prefers to procrastinate. Knowing this, self $T-2$ prefers to perform the task immediately because he prefers it to doing it at period T. The rest of the argument is straightforward. The conclusion is that in subgame perfect equilibrium, the decision maker performs the task at period k if and only if $T-k$ is even. This result has the flavor of the well-known finite horizon paradoxes, which result from the unintuitiveness of applying backward induction reasoning for a large number of iterations. As such, it demonstrates the limitations of the multi-selves approach discussed earlier.

By comparison, a fully naive decision maker (i.e., a decision maker who believes at every t that $\succ_s = \succ_t$ for every $s > t$) will procrastinate until the final period. The reason is simple: at every period $t < T$, he believes that at $t + 1$ he will decide to perform the task right away, which he never does, until he is forced to perform is at period T. Note that this result is also puzzling, in the sense that the fully naive decision maker never learns anything about his true propensity to procrastinate. A more realistic model would allow the decision maker to learn in the course of the dynamic decision problem that he has a tendency to procrastinate.

2.3 Welfare Analysis in Multi-Selves Models

Once we perceive the decision maker as a collection of different selves with different preferences over the set of consequences, and analyze the decision maker's behavior as the outcome of a game-theoretic solution concept in the game played among the selves, welfare evaluation becomes problematic. Should we view the different selves as genuinely separate entities, and conduct welfare analysis in the same way that we do in standard, inter-personal games? If so, what would constitute a proper "social" welfare function? Will our intuitions regarding this question be the same intuitions that we apply to standard, inter-personal situations?

In some cases, even the primitives of such a quasi-social-choice analysis are ill-defined. For instance, in infinite-horizon saving-consumption models, our application of the multi-selves approach requires us to write down each self's preferences over *continuation* paths. In order to view the collection of selves as a society in the sense of social choice theory, we need to define their preferences over the same domain (i.e., complete consumption paths $(c_1, c_2, ...)$). Is there a natural completion of the preferences of the self that acts at some period $t > 1$ to this domain? Is it legitimate to define preferences over things that these selves cannot choose by definition?

In many applications of the multi-selves approach, the preferences of some of the selves are viewed as *vices*. For instance, when the model describes a decision maker's attempt to maintain a diet, then first-period considerations appear to have a higher normative status than second-period considerations, because the former reflect long-run planning whereas the latter reflect visceral urges. In these examples, the decision maker's first-period commitment rankings are assumed to represent his welfare. In the saving-consumption model described above, many economists tend to view the taste for immediate gratification, captured by the parameter β, as a vice. Welfare analysis then proceeds according to preferences over consumption paths that are represented by the same formula, using $\beta = 1$. Note that in this case, the welfare criterion does not coincide with the preferences of any of the selves.

Of course, the judgment of what constitutes a "genuine" preference and what constitutes an urge or a vice is highly subjective. For instance, imagine that you have just bought Kafka's "The Trial". You intend to read the first few pages and then go to bed early, because you are planning to get early for day job (you're an accountant). However, as you read the book you become engrossed and finish reading it at 4AM, leaving only one hour of sleep. Is it obvious now what is the genuine preference and what is the vice? Would our answer change if we substitute a porn movie for the literary masterpiece? Would it affect our judgment if we knew that you felt regret or satisfaction in the aftermath of this experience (i.e., in a "third" period which is behaviorally irrelevant for the decision problem)?

These examples demonstrate that when we conduct normative analysis in multi-selves models, we need to define a notion of the decision maker's welfare, quite separately from the preferences we assign to the different selves. And we need to justify these assumptions just as we need to justify behavioral assumptions in positive analysis. Developing a specialized "social choice theory" that would organize our intuitions about welfare analysis in multi-selves models is an interesting topic for future research.

3 Interaction between a Profit-Maximizing Firm and a Dynamically Inconsistent Consumer

In this section we apply the tools presented in this lecture to situations in which a monopolistic firm interacts with a consumer having changing tastes. The consumer may be naive or sophisticated. Consider the following simple two-period model. Let $X = [0, 1]$ be the set of feasible quantities of a product or service provided by a monopolistic firm. A consumer can choose a quantity in period 2, conditional on having signed a contract with the firm in period 1. The consumer has quasi-linear preferences. However, these change over time. In period 1, the consumer's preferences over quantity-payment pairs (x, t) are represented by the utility function $u(x) - t$, whereas in period 2, they are represented by the utility function $v(x) - t$. In period 1, the consumer is indifferent between the quantity-payment pair $(0, 0)$ and not signing any contract. Assume that the firm has a constant marginal cost of c.

Throughout this section, we will assume that $u(0) = v(0) = 0$, and that both u and v are twice differentiable, with $u', v' > 0$, $c < u'(0) < \infty$, $c < v'(0) < \infty$, $u'', v'' < 0$, $u'(1), v'(1) < c$. We will examine two cases: *(i)* $u'(x) < v'(x)$ for all x; *(ii)* $u'(x) > v'(x)$

for all x. The first case fits situations in which the consumer's willingness to pay for any additional unit of the product or service increases as a result of the passage of time, or as a result of some event associated with the act of signing a contract with the firm. For instance, think of x as the amount of credit extended to the consumer by a credit company, or as the number of minutes the consumer talks on the mobile phone. The second case fits situations in which consuming the product or service is rewarding in the long run but taxing in the short run (exercising in a gym, consuming "high" culture, learning to play a musical instrument, etc.).

What is the structure of contracts that maximize the firm's profit? Does it account for features of contracts that we see in reality? What is the role of consumer naivety in determining the structure of optimal contracts? Can a firm offer a menu of contracts that will screen the consumer's "type" (naive or sophisticated)?

Comment. We assume that the firm assigns probability one to the consumer's second-period preferences being given by v, whereas a fully naive consumer assigns probability one to his preferences being given by u. Technically, this results in a game-theoretic interaction between two players with non-common priors. We, the modelers, take sides in this disagreement and accept the firm's opinion, thus referring to the consumer's beliefs as biased or naive. However, this is only one possible interpretation: we could refrain from taking sides and assigning a truth value to the parties' beliefs. Whatever interpretation we choose, what is important is that neither party thinks that he lacks some piece of information that the other party possesses. In particular, our consumer would disagree with our judgment that his beliefs about second-period preferences are "naive".

Comment. A contract offered to a naive consumer is essentially a bet on his second-period preferences, where the bet is motivated by the non-common priors. If the two parties could directly condition on the consumer's second-period preferences, it would be beneficial to both to agree on infinite bets. However, since they can only condition on the consumer's action, and since both u and v are bounded, the stakes of the bet are bounded as well.

3.1 Two-Part Tariffs

Let us first restrict attention to two-part tariffs, namely contracts of the form $t(x) = 0$ if $x = 0$ and $t(x) = A + px$ for $x > 0$. The restriction to two-part tariffs is made for several reasons. First, they are common. Second, two-part tariffs are known to be optimal when the consumer is dynamically consistent and his preferences satisfy the condition imposed on u and v. Finally, while optimal two-part tariffs imply $p^* = c$ in this case, we will see that when the consumer is dynamically inconsistent, the model can account for real-life departures from marginal-cost pricing.

Suppose that the consumer is known to be sophisticated. Then, he knows that given a two-part tariff (A, p), he will choose in period 2 the quantity x^v defined by $v'(x^v) = p$. The maximal A that he is therefore willing to accept in period 1 satisfies $A = u(x^v) - px^v$. Thus, the firm effectively chooses x^v to maximize $px^v + [u(x^v) - px^v] - cx^v = u(x^v) - cx^v$, yielding $u'(x^v) = c$. But we have seem that $v'(x^v) = p$. It follows that in case *(i)*, where $v' > u'$, we obtain $p^* > c$, and in case *(ii)*, where $v' < u'$, we obtain $p^* < c$. That is, the optimal two-part tariff departs from marginal-cost pricing. When the consumer's marginal willingness to pay increases in period 2, the price per unit is above marginal cost. And when the consumer's marginal willingness to pay diminishes in period 2, the price per unit is below marginal cost. In addition, note that the optimal A need not be positive. If the gap between first- and second-period willingness to pay is sufficiently large, A^* is negative.

Now suppose that the consumer is known to be fully naive - i.e., he expects his second-period preferences to be given by u rather than v. The firm expects the consumer to choose $x^v(p) = \arg\max[v(x) - px]$. However, in period 1 the consumer expects that he will choose $x^u(p) = \arg\max[u(x) - px]$. The maximal A that he is therefore willing to accept in period 1 satisfies $A = u(x^u) - px^u$. Thus, the firm effectively chooses p to maximize $px^v(p) + [u(x^u(p)) - px^u(p)] - cx^v(p)$. Note that if the firm sets $p = c$, the objective function is reduced to $u(x^u(p)) - px^u(p)$, which is the same as in the case of a sophisticated consumer. Thus, the firm's profit from the interaction with a naive consumer is at least as high as its profit from the interaction with a sophisticated consumer.

Let us rewrite the firm's maximization problem:

$$\max_{p} \quad [u(x^u(p)) - px^u(p)] + (p - c)x^v(p)$$

subject to

$$\begin{aligned} x^u(p) &= \arg\max u(x) - px \\ x^v(p) &= \arg\max v(x) - px \end{aligned}$$

The derivative of the objective function with respect to p is $x^v(p) - x^u(p) + \frac{dx^v(p)}{dp} \cdot (p - c)$.

In case *(i)*, since $x^v(p) > x^u(p)$, and since $\frac{dx^v(p)}{dp} < 0$, this means that the derivative is positive for every $p \leq c$. Hence, it is optimal for the firm to set $p^* > c$. In case *(ii)*, since $x^v(p) < x^u(p)$, this means that the derivative is negative for every $p \geq c$. Hence, it is optimal for the firm to set $p^* < c$.

Thus, both when the consumer is sophisticated and when he is naive, the optimal two-part tariff induces the same qualitative departure from marginal-cost pricing. When the consumer is sophisticated, this is meant as a commitment device, whereas when he is naive, it is meant as a device for exploiting his naivety. This finding is consistent with observed pricing behavior in several industries that fit these cases. Credit card companies often offer "freebees" in order to tempt consumers to use credit cards, and at the same time charge exceedingly high interest rates for the credit they extend. At the other extreme, health clubs charge high one-off subscription fees in conjunction with low per-visit payments. Note that the departure from marginal-cost pricing is qualitatively the same whether the consumer is sophisticated or naive. This raises the question of whether two-part tariffs retain their optimality when we turn from dynamically consistent to dynamically inconsistent consumers. Let us address this problem now.

3.2 Optimal Contract Design

When designing the optimal contract for a consumer who is known to be sophisticated, we can assume that without loss of generality, the firm needs to fix a pair (x^v, t^v), where x^v is the quantity the consumer chooses in period 2 and t^v is the amount that he pays in period 2, since the two parties agree that the consumer's second-period preferences will be given by v. We can set $t(x)$ to be arbitrarily large for any $x \neq x^v$, without loss of generality. The firm's maximization problem is the reduced to:

$$\max_{x^v, t^v} \quad t^v - cx^v$$

subject to

$$u(x^v) - t^v \geq 0$$

which means that the firm effectively chooses x^v to maximize $u(x) - cx$. Thus, as far as the first-period self is concerned, the contract induces an efficient action, and the firm manages to extract the self's surplus. The reason is simple: since the firm is able to offer a perfect commitment device ($t(x)$ being arbitrarily large for any $x \neq x^v$), it can induce the outcome that is considered efficient from the first self's point of view. Note that this conclusion is independent of any assumption on u. However, our assumptions on u imply that the maximal profit can be attained with a two-part tariff.

Things are different when we turn to a fully naive consumer. Here, the firm needs to fix a 4-tuple, (x^u, t^u, x^v, t^v), where x^ω is the quantity the consumer chooses in state ω and t^ω is the amount he pays in state ω. We can set $t(x)$ to be arbitrarily large for any $x \neq x^u, x^v$, without loss of generality. The firm's maximization problem is:

$$\max_{x^u, t^u, x^v, t^v} \quad t^v - cx^v$$

subject to

$$v(x^v) - t^v \geq v(x^u) - t^u \tag{3}$$

$$u(x^u) - t^u \geq u(x^v) - t^v \tag{4}$$

$$u(x^u) - t^u \geq 0 \tag{5}$$

The first constraint ensures that the consumer indeed chooses x^v when his preferences are given by v, as the firm expects. The second constraint ensures that he chooses x^u when his preferences are given by u, as the consumer expects. The final constraint is a standard participation constraint.

It is easy to see that constraint VR is binding in optimum. Thus, $t^v = t^u + v(x^v) - v(x^u)$. But this means that IR is also binding in optimum. It follows that $t^v = u(x^u) - v(x^u) + v(x^v)$. Thus, the firm's maximization problem is reduced to finding x^u and x^v which maximize $v(x^v) - cx^v + u(x^u) - v(x^u)$. Let us show that if the optimal contract yields higher profits than a perfect commitment device, then we can set $x^u = \arg\max[u(x) - v(x)] \equiv x^*$ without upsetting any of the constraints. Assume the contrary. If $x^u = x^v$, the contract is a perfect commitment device. Suppose that $x^u \neq x^v, x^*$ and imagine that the firm deviates from x^u to x^* and modifies t^u into t^*, defined by $u(x^*) - t^* = u(x^u) - t^u$. Then, constraints UR and IR obviously remain unchanged. And by the definition of x^*, it is easy to see that constraint VR continues to hold.

It follows that in optimum, we can set $x^u = \arg\max[u(x) - v(x)]$ and $x^v = \arg\max[v(x) - cx]$, so that $t^v = v(x^v) + u(x^*) - v(x^*)$ and $t^u = u(x^*)$. (Note that if $x^u = x^v$, we obtain the optimal contract for the sophisticated consumer.) Note that in case *(i)*, $x^* = 0$ and in case *(ii)*, $x^* = 1$. That is, when the consumer's actual willingness to pay increases in period 2, the firm designs the contract so that the naive consumer believes that he will not buy any quantity. In contrast, when the consumer's actual willingness to pay decreases in period 2, the firm designs the contract so that the naive consumer believes that he will purchase the maximal quantity. The quantity chosen by the consumer according to the firm's expectation is efficient from the second-period self's point of view. Given out assumptions on u and v, this means that $v'(x^v) = c$. This is different from the value induced by the optimal two-part tariff. With a two-part tariff, the only way to get $v'(x^v) = c$ is to set $p = c$. But as we saw earlier, this price does not allow the firm to attain a profit which is higher than what it can get when facing a sophisticated consumer. Thus, the restriction to two-part tariffs, while carrying no loss of generality in the case of sophisticated consumers, is sub-optimal in the case of naifs.

3.3 Screening the Consumer's Type

If the firm does not know whether it is facing a sophisticated or a naive consumer, it can nonetheless screen his type at no cost, simply by offering a menu of contracts consisting of the optimal contract for each type. Both consumer types will agree that the contract designed for a sophisticated consumer provides a perfect commitment device, and therefore both consumer types will evaluate it at 0. However, the two types differ in their evaluation of the contract designed for the naive consumer. While the naive consumer evaluates it at 0 (because the participation constraint is binding), the sophisticated consumer believes that it is an exploitative contract which yields a negative payoff (this is by definition, since as we saw, it generates a higher profit for the firm than any other contract a sophisticated consumer would accept), and therefore he will not choose it. it follows that a menu consisting of these two contracts will achieve perfect screening.

This perfect-screening result relies on the assumption that there are two types, one of which being fully sophisticated. Suppose that in addition to these types, there are also types who are partially naive, in the sense that they assign probability α to u and probability $1 - \alpha$ to v, where $\alpha \in (0, 1)$. In that case, screening the consumer's degree of naivety would entail a cost. The reason is that as long as a consumer is not fully sophisticated, the firm would want to offer him a contract which is not a perfect commitment device, so that $x^u \neq x^v$, and such a contract would be evaluated differently by different types. In particular, a naive consumer type may have an incentive to choose the contract aimed at at a more sophisticated consumer type.

The optimal menu of contracts has an interesting welfare property. While the contract aimed at the sophisticated consumer induces an outcome which is ex-post efficient according to u, the contract aimed at the naive consumer induces an outcome which is ex-post efficient according to v.

Comment: The interpretation of naivety in an interactive context. We have already commented on the question of how to interpret the full-naivety assumption in the consumer-firm model. One interpretation is that the consumer is simply unaware of the possibility that his tastes will change. Another interpretation is that while he is aware of this possibility, as well as the fact that the firm thinks differently, he is confident of his own belief and "agrees to disagree" with the firm. Given the structure of the optimal contract aimed at the naive consumer, we should revisit this question. Consider case *(i)*, in which v lies above u. We saw that in this case, the optimal contract is such that $x^* = 0$ - i.e., the naive consumer believes that he will buy zero units and pay nothing. Even if the consumer is unaware of the possibility that his tastes will change, he should ask himself why the firm would go through the trouble of offering him a contract which induces zero consumption. And one conclusion might be that the firm believes that his tastes will change. Therefore, even if the consumer begins the interaction in a state of unawareness, the structure of the optimal contract might trigger a chain of reasoning that will bring the possibility of changing tastes to his awareness.

This problem becomes more acute when the firm offers a menu of contracts, in order to screen the consumer's type. For some configurations of u, v, the optimal menu has an interesting property - the payment associated with each outcome is always (weakly) higher under the contract aimed at the sophisticated consumer, and for some outcomes the payment is strictly higher. In other words, the naive consumer will find that the other contract in the menu is *dominated* by the contract he chooses.[1] Moreover, he may observe that other consumers, who same to share his current tastes, choose the other contract. Why would they choose a dominated contract, unless they shared the firm's belief that their tastes might change? This thought may lead the naive increase to question his beliefs.

Thus, we see that although the naivety assumption may be interpreted in a variety of ways, the structure of the optimal menu of contracts may render some of these interpretations implausible. And this may also be a consideration that the monopolist takes into account when devising a menu of contracts. Although he wants to exploit the naive consumers, he may try to avoid making the exploitation too transparent, so that the naive consumers will not be triggered to reconsider their beliefs.

[1] For instance, let $v(x) = x$ and let $u(x) = 2x$ for $x \in [0, \frac{1}{2}]$ and $u(x) = 2(1-x)$ for $x \in (\frac{1}{2}, 1]$. Showing the dominance property of the optimal menu is left for the reader as an exercise.

4 Bibliographic Notes

For works on intertemporal choice patterns and their interpretation, see Frederick et al. (2002) and Rubinstein (2003). The multi-selves approach was originally introduced by Strotz (1956) and developed further by Peleg and Yaari (1973). The axiomatic characterization of first-period behavior induced by the multi-selves approach is a variant on a result due to Gul and Pesendorfer (2001). The procrastination example is based on O'Donoghue and Rabin (1999). The saving-consumption model is based on Phelps and Pollack (1968) and Laibson (1997). The firm-consumer application is based on DellaVigna and Malmendier (2004) and Eliaz and Spiegler (2006).

5 Exercises

1. (*Based on Rubinstein (2006a)*) Consider the preferences of a decision maker over infinite sequences of monetary payoffs $(x_t)_{t=1,2,...}$. The preferences are represented by a standard, exponential discounting utility function $\Sigma_{t=1,2,...}\delta^{t-1}x_t$. Suppose that if $x_k = 1$, $x_{k+1} = 0$, $y_k = 0$, $y_{k+1} = 2$ and $x_t = y_t$ for every $t \neq k, k+1$, then $x \succ y$. Show that for every $m \geq 1$, the sequence $(1, 0, 0, ...)$ is preferred to the sequence: $x_k = 0$ for every $k \leq m$ and every $k > 2^m + m$, and $x_k = 1$ for every $m < k \leq 2^m + m$. Try to interpret this result and link it to the calibration theorem due to Rabin (2000).

2. (*Based on Rubinstein (2006b)*) At every period $t = 1, ..., T$, a consumer is offered a product labeled G with probability $\frac{1}{2}$ or a product labeled B with probability $\frac{1}{2}$. If the consumer accepts the offer, the decision problem is terminated. If he declines it at $t < T$, he moves to the next period. If the consumer declines all first $T-1$ offers, he is forced to accept any offer at period T. An outcome in the consumer's decision problem is a pair (x, t), where $x \in \{G, B\}$ and $t \in \{1, ..., T\}$. The consumer has (β, δ) preferences over outcomes, where $\delta = 1$. Let g and b denote the instantaneous payoff from G and B, respectively, and assume that $g > \beta(g+1) > b > \beta g$. A social planner prefers that the consumer ultimately chooses G rather than B. However, she cannot verify this choice. For a fixed termination period, this is also the consumer's preference. However, the consumer is impulsive and prefers accepting B at period t to waiting for future offers. The social planner has a budget of \$1, and is able to commit in advance to a subsidy scheme: a payment to the consumer which is conditioned on the termination period. Assume that the consumer is naive. Show that as T tends to infinity,

the social planner can devise a subsidy scheme for which the probability that the consumer will choose G is almost one.

References

[1] DellaVigna, S. and U. Malmendier (2004), “Contract Design and Self-Control: Theory and Evidence”, *Quarterly Journal of Economics* **119**, 353-402.

[2] Eliaz, K. and R. Spiegler (2006), “Contracting with Diversely Naive Agents”, *Review of Economic Studies* **73**, 689-714.

[3] Frederick, S., G. Loewenstein and T. O’donoghue (2002), “Time Discounting and Time Preference: A Critical Review”, *Journal of Economic Literature* **40**, 350-401.

[4] Gul, F. and W. Pesendorfer (2001), “Temptation and Self-Control”, *Econometrica* **69**, 1403-1435.

[5] O’Donoghue, T. and M. Rabin (1999), “Doing It Now or Later”, *American Economic Review* **89**, 103-124.

[6] Peleg, B. and M. Yaari (1973), “On the Existence of a Consistent Course of Action when Tastes are Changing”, *Review of Economic Studies* **40**, 391-401.

[7] Rabin, M. (2000),. “Risk Aversion and Expected-Utility Theory: A Calibration Theorem”, *Econometrica* **68**, 1281-1292.

[8] Rubinstein A. (2003), “Economics and Psychology? The Case of Hyperbolic Discounting”, *International Economic Review* **44**, 1207-1216.

[9] Rubinstein A. (2006a), “Dilemmas of An Economic Theorist”, *Econometrica* **74**, 865-883.

[10] Rubinstein A. (2006b), “Comments on Behavioral Economics”, forthcoming in *Adavnces in Economic Theory* (World Congress of the Econometric Society).

[11] Strotz R. (1956), “Myopia and Inconsistency in Dynamic Utility Maximization”, *Review of Economic Studies* **23**, 165-180.