Chapter 3
Dynamically Inconsistent Preferences II: Imperfect Monitoring

Ran Spiegler

December 4, 2008

In the previous chapter, we introduced the multi-selves model of dynamically inconsistent consumers, and used it to study the pricing behavior of firms who face such consumers. We assumed perfect monitoring, in the sense that firms were perfectly able to condition payments on consumers' second-period actions. For sophisticated consumers, perfect monitoring is important because it ensure the feasibility of perfect commitment devices such consumers seek. For naive consumers, perfect monitoring is important because it allows the firm to devise an ornate "betting" contract.

In this section we relax the perfect monitoring assumption in various dimensions. First, we restrict attention to two-part tariffs, capturing situations in which firms are unable to discriminate on the basis of consumed quantity. Next, we will see what happens when firms are unable to prevent consumers from renegotiating the terms of the contract or buying from other suppliers in the second period. Finally, we will explore the role of consumers' own self-control in situations where perfect commitment devices are infeasible. For this purpose, we will need to extend the multi-selves model, which itself is unable to capture self-control. In all three cases, we will see that the distinction between perfect and imperfect monitoring is important in the context of interactions between profit-maximizing firms and consumers with changing tastes, and leads to new insights.

1 Two-Part Tariffs

In this section we assume that $X = [0, 1]$ and restrict attention to pricing schemes that take the form of two-part tariffs: $t(x) = 0$ 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 commonly observed. Second, the restriction is known to be without loss of generality when consumers have dynamically consistent, quasi-linear preferences that satisfy diminishing marginal willingness to pay. Finally, while optimal two-part tariffs in standard models have the property that the per-unit price p is equal to marginal cost, we will see that when the consumer is dynamically inconsistent, the model can account for real-life departures from marginal-cost pricing.

Throughout this section, we assume that $u(0) = v(0) = 0$, and that both u and v are continuously twice differentiable, with $u', v' > 0$ and $u'', v'' < 0$. We also assume that the firm faces a constant marginal cost κ, where $\kappa < u'(0) < \infty$, $\kappa < v'(0) < \infty$, $u'(1), v'(1) < \kappa$. 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 it is tempting to consume more in period 2 (e.g. when x denotes the amount of food consumption and the consumer is on a diet). The second case fits situations in which consuming the product is rewarding in the long run but taxing in the short run (e.g. when x denotes the intensity of physical exercise, which is beneficial in the long run but taxing in the short run).

1.1 Departure from Marginal Cost Pricing

We will restrict attention to monopoly pricing.

Proposition 1 *When the consumer is known to be fully sophisticated, the monopolist will choose a two-part tariff* (A^*, p^*) *such that in case (i),* $p^* > \kappa$ *and in case (ii),* $p^* < \kappa$.

Proof. 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 level of consumption 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 p to maximize $\pi_s(p) = px^v(p) + [u(x^v(p)) - px^v(p)] - \kappa x^v(p) = u(x^v(p)) - \kappa x^v(p)$, where $x^v(p)$ is given by $v'(x) = p$, yielding $u'(x^v) = \kappa$. But we have seem that $v'(x^v) = p$. It follows that in case *(i)*, where $v' > u'$, we obtain $p^* > \kappa$, and in case *(ii)*, where $v' < u'$, we obtain $p^* < \kappa$. 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. ■

Note that in case *(i)*, the optimal value of A may be *negative*. This captures immediate benefits that consumers receive when they sign up for the service. In case *(ii)*, the optimal value of A is necessarily positive.

It turns out that these qualitative conclusions hold also when the consumer is known to be fully naive.

Proposition 2 *When the consumer is known to be fully naive, the monopolist will choose a two-part tariff* (A^*, p^*) *such that in case (i),* $p^* > \kappa$ *and in case (ii),* $p^* < \kappa$.

Proof. Suppose that the consumer is known to be fully naive - i.e., he expects his second-period preferences to be given by u. 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 $A = u(x^u(p)) - px^u(p)$. It follows that the firm effectively chooses p to maximize $u(x^u(p)) + p \cdot [x^v(p) - x^u(p)] - \kappa x^v(p)$. For expositional simplicity, we will analyze the problem as if both $x^u(p)$ and $x^v(p)$ are always strictly between 0 and 1.

In the previous chapter, we observed that the monopolist's profit is higher when the consumer is naive than when he is sophisticated. This conclusion extends to the restricted domain of two-part tariffs. To see why, consider a two-part tariff (A, p) that leaves a sophisticated consumer indifferent between accepting and rejecting the contract. The firm's profit from such a two-part tariff is equal to $u(x^v(p)) - \kappa x^v(p)$. Since u is concave and $x^u(p)$ is given by $u'(x) = p$:

$$u(x^u(p)) + p \cdot [x^v(p) - x^u(p)] > u(x^v(p)) \qquad (1)$$

Thus, the same tariff (A, p) would generate a strictly higher profit if the consumer were naive. It follows that the firm's maximal profit when it faces naifs is strictly higher than when it faces sophisticates.

Let us rewrite the firm's maximization problem when it faces a naive consumer:

$$\max_p \; [u(x^u(p)) - px^u(p)] + (p - \kappa)x^v(p)$$

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 $x^v(p) - x^u(p) + \frac{dx^v(p)}{dp} \cdot (p - \kappa) + \frac{dx^u(p)}{dp} \cdot (\frac{du}{dx}(x^u(p) - p)$

Let us evaluate this derivative in case *(i)*. By our assumption that $v' > u'$, $x^v(p) - x^u(p) > 0$. Since $x^u(p)$ is given by $u'(x) = p$ and $v' > u'$, we have $u'(x^u) - p < 0$. By the concavity of u and v, $\frac{dx^v(p)}{dp}$ and $\frac{dx^u(p)}{dp}$ are negative. It follows that the overall derivative of the objective function with respect to p is positive for every $p \leq \kappa$. Hence, it is optimal for the firm to set $p^* > \kappa$. In case *(ii)*, a similar calculation implies that it is optimal for the firm to set $p^* < \kappa$. ■

Therefore, 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.

A real-life example: health clubs

An empirical study by DellaVigna and Malmendier (2006) examined consumer behavior in the context of health clubs. Working out regularly at a health club is an activity that is costly in the short run and rewarding in the long run. Therefore, consumers with a taste for immediate gratification will display dynamically inconsistent preferences. From an ex-ante point of view, they would like to commit to go to the gym on a regular basis. However, when the time comes to carry out this commitment, consumers may turn lazy and procrastinate. Health clubs generally offer two types of payment schemes: membership and non-membership. Roughly speaking, this is a menu consisting of two kinds of two-part tariffs. A member makes a large up-front payment for a long period, and subsequently faces a small or no per-visit price. A non-member simply pays a relatively high per-visit price. DellaVigna and Malmendier (2006) gathered data on consumers' contract choice and subsequent attendance. They showed that many of the consumers who chose the membership option attended the gym so irregularly that they would have been better off if they had stuck to the non-membership option. That is, the effective per-visit price implicit in their total payment and actual attendance is higher than the explicit per-visit price under the non-membership plan. A natural interpretation of this effect is that some of the consumers who opt for the membership plan are naive. They falsely believe that their current strong willingness to workout will persist in the future, and therefore find the membership plan more attractive than they would find it if they correctly anticipated their future preferences.

1.2 Welfare Analysis

Is the restriction to two-part tariffs an effective constraint for the monopolist? We saw that when the consumer is sophisticated, the optimal two-part tariff reproduces the same second-period action as the optimal unrestricted contract, and fully extracts the consumer's first-period willingness to pay for this action. Thus, the restriction to two-part tariffs is without loss of generality. (Of course, this conclusion is relevant only for the environments we restricted attention to, in terms of assumptions on u and v, in which two-part tariffs are of interest in the first place.)

In contrast, the restriction to two-part tariffs carries a loss of generality in case of a naive consumer. Recall that the optimal unrestricted contract induces an action x^v which is efficient according to v. Given our assumptions on v, this means $v'(x^v) = \kappa$. But the optimal two-part tariff induces an action x^v that maximizes $v(x) - p^* x$, and since $p^* \neq \kappa$, the outcome is inefficient according to v. Hence, when facing a naive consumer, the monopolist could do better by offering a pricing scheme which does not fall into the form of a two-part tariff.

The intuition for this difference is as follows. When dealing with a naive consumer, the monopolist is interested in two actions: the action x^v it expects the consumer to take and the action x^u the consumer expects to take. The restriction to two-part tariffs implies that x^v and x^u are chosen such that $v'(x^v) = u'(x^u)$. In contrast, the optimal unrestricted contract induces $x^v \in \arg\max[v(x) - \kappa x]$ and $x^u \in \arg\max[u(x) - v(x)]$, and there is no particular reason these solutions should coincide. In contrast, when dealing with a sophisticated consumer, the monopolist is interested in only one outcome: x^v. This is also the outcome the consumer expects to happen. Therefore, by setting $p = v'(x^v)$, the monopolist can implement the optimal commitment device.

Let us now turn to the implications of the restriction to two-part tariffs for the welfare of naive consumers. Under the optimal two-part tariff (A, p), the first self's net payoff is $u(x^v_{tpt}) - [u(x^u_{tpt}) + p \cdot (x^v_{tpt} - x^u_{tpt})]$, where x^v_{tpt} and x^u_{tpt} are given by $v'(x^v_{tpt}) = u'(x^u_{tpt}) = p$, respectively. As we saw in the proof of Proposition 2 - see in particular inequality (1) - this net payoff is negative. It follows that the restriction to two-part tariffs does not eliminate the exploitation of naive consumers (from the point of view of their first-period preferences).

Does the restriction to two-part tariffs at least curb exploitation, relative to the optimal unrestricted contract? Let us examine case (i), in which $v' > u'$. Recall that under the optimal unrestricted contract derived in Chapter 2, the first-period self's net payoff is $u(x^v_{opt}) - v(x^v_{opt})$, where x^v_{opt} is given by $v'(x) = \kappa$. Recall also that under case *(i)*, $p > \kappa$. Therefore, $x^v_{opt} > x^v_{tpt}$. By our assumptions on u and v,

$u(x^v_{opt}) - v(x^v_{opt}) < u(x^v_{tpt}) - v(x^v_{tpt}) < 0$. Since x^u_{tpt} is given by $u'(x^u_{tpt}) = p$. In case *(i)*, this implies $u(x^u_{tpt}) + p \cdot (x^v_{tpt} - x^u_{tpt}) < v(x^v_{tpt})$. Combining these inequalities, we obtain that the first-period self's net payoff is strictly lower under the optimal contract relative to the optimal two-part tariff. In other words, the restriction to two-part tariffs mitigates the exploitation of naive consumers, even if it does not eliminate it altogether.

Exercise 1 *Show that the same welfare analysis holds under case* (*ii*), *i.e.* $v' < u'$. *(Hint: Use a diagrammatic characterization of* $x^v_{opt}, x^u_{opt}, x^v_{tpt}, x^u_{tpt}$*.)*

This finding is a first illustration in this book of the following principle: optimal pricing schemes for boundedly rational consumers often contain a complexity that would not be used if the consumer were rational. Therefore, forcing the firm to forego this complexity in favor of a simple pricing scheme may mitigate the exploitation of the boundedly rational consumer.

2 Destabilization of Commitment Devices: Renegotiation and Spot Market Competition

In this sub-section we introduce the possibility that consumers who signed a contract in period 1 will face new alternatives in period 2. This may take several forms. First, the same firms that offered the first-period contract may try to renegotiate. Second, other firms, which did not sign contracts with the consumers in period 1, may enter the market and offer "spot" contracts that compete with the contracts the consumers accepted in period 1.

Renegotiation

Second-period renegotiation can destabilize the commitment device chosen by sophisticated consumers in period 1. Recall that a competitive contract for sophisticates, denoted t_s, induces an action x^v_s that maximizes $u - c$ and a payment $t^v_s = c(x^v_s)$, and thus generates zero profits for the firms. Typically, x^v_s does not maximize $v - c$, and so a firm can propose replacing the original contract with a new contract t_r, which induces an action x^v_r that maximizes $v - c$ and a payment $t^v_r > c(x^v_r)$, such that $v(x^v_r) - t^v_r > v(x^v_s) - t^v_s$. Both the firm and the sophisticated consumer will prefer the new contract in period 2.

Thus, the only commitment devices signed in period 1 that are stable with respect to renegotiation are those that induce an action $x^v \in \arg\max(v - c)$ and a payment $t^v = c(x^v)$. Such commitment devices enforce an outcome which is efficient according to the second-period self's preferences, but inefficient according to the first-period self's preferences. Thus, the possibility of renegotiation distorts the interaction with sophisticated consumers.

The possibility of renegotiating commitment devices is interesting from a legal point of view. According to a powerful tradition in legal theory, when two parties agree to renegotiate an existing contract, the court should not void the newly signed contract. However, the rationale behind this libertarian stance is typically that the renegotiation was a result of new information, whereas in our case, the renegotiation is a result of predictably changing tastes. The ability to enforce commitment contracts thus calls for a legal doctrine that acknowledges the distinction between these different motivations for renegotiation.

Spot markets

Second-period competition from "spot" contracts is another force that can destabilize commitment contracts signed in period 1. Suppose that the commitment device sophisticates choose in period 1 induces a low level of consumption, and that second-period preferences exhibit *increased* willingness to pay for additional units. Then, the consumer can obtain the extra units from another supplier without breaking the first-period contract. For instance, think of a consumer who signs in period 1 a loan contract with his bank, in the hope that the stringent terms of the contract will deter him from over-borrowing in period 2. Spot market competition in period 2 can come from other lenders, such as credit card companies and loan sharks. The first-period equilibrium commitment device which is immune to such spot market competition is the same renegotiation-proof contract that was described above.

Renegotiation and spot market competition are irrelevant for the contracts aimed at naive consumers, because these contracts are efficient according to the second-period self's preferences. If the first-period market is monopolistic, then spot market competition in period 2 can certainly lower prices, but if the first-period market is competitive, spot market competition in the second period will have no effect on naifs.

Exercise 2 *Recall the example in Chapter 2 of a competitive model with diversely naive consumers, under the probabilistic definition of partial naivety. Redo the analysis under the assumption that in period 2 there is a competitive market for spot contracts.*

2.1 Addictive Goods

To be added later.

3 Self-Control

When consumers do not have at their disposal a perfect commitment device, an alternative channel for overcoming temptations is the exertion of self-control. A consumer with self-control acknowledges the temptation yet manages to implement his first-period preferences to some degree. In this section I present an extension of the multi-selves model that incorporates self-control in a particular manner, and analyze some of its implications for firms' pricing behavior.

The cornerstone of the model of self-control is that consumer preferences are defined not over the set of consumption decisions Z, but over the extended set of *decision paths* $\{(A, z) \mid \varnothing \subset A \subseteq Z,\ z \in A\}$. I will often refer to choice sets A as "*menus*".

The idea of preferences that depend on the entire decision path - in particular, on unchosen elements in the choice set - is quite appealing as a way to model self-control. When a consumer who exercises self-control evaluates his consumption decision, he takes into account not only the action he eventually took, but also feasible actions he was tempted to take yet chose not to by exerting self-control. The fact that he has to exert self-control in order to avoid tempting actions means that he is not indifferent to their inclusion in the menu.

In the model we present, exerting self-control is costly. The greater the temptations in the choice set, the higher the cost of self-control. This cost may be purely mental, or it can be the cost of a physical activity which is not explicitly modeled as an element in Z (e.g., participation in a support group). The following class of utility functions that represent preferences over decision paths captures this consideration:

$$U(A, z) = u^*(z) - s(A, z)$$

where $u^*(z)$ is the "commitment utility" attached to z - namely, the payoff that z yields when it is the only available element, and $s(A, z)$ is the *cost* of the self-control needed to choose z from A.

We will employ a particular specification of the cost of self-control function:

$$s(A, z) = \max_{y \in A} v^*(y) - v^*(z)$$

where v^* is the "temptation utility" attached to z. The self-control cost of choosing z from A is simply the difference in temptation utility between z and the maximally tempting element in A. This leads to an alternative way of writing the utility function over decision paths:

$$U(A, z) = [u^*(z) + v^*(z)] - \max_{y \in A} v^*(y)$$

Note that since the last term is fixed given a menu A, the consumer's second–period choices are consistent with maximizing the utility function $u^* + v^*$ defined over Z. Thus, second-period choices reflect a compromise between "commitment preferences" and "temptation preferences".

If we only observed how the consumer chooses from menus in period 2 and how he ranks singleton menus in period 1, we would find this observed behavior consistent with a multi-selves model in which the first-period self has preferences over Z represented by u^* and the second-period self has preferences over Z represented by v^*. This inconsistency is rationalized by a consistent preference relation over the extended domain of decision paths. The first-period ranking of singleton menus is viewed as a preference over decision paths of the form $(\{x\}, x)$, while choices from a menu A in period 2 are interpreted as preferences over decision paths of the form (A, x).

Thus, while the multi-selves model views first-period and second-period choices as reflecting inconsistent preferences over Z, the self-control model views them as consistent preferences over an expanded consequence space, thus subsuming the inconsistent preferences into a consistent ranking over extended consequences. This apparent elimination of dynamic inconsistency has no behavioral content: it is purely a matter of language. One can always make inconsistent choices appear consistent by claiming that they are outputs of different choice problems.

This is not to say that the self-control model is behaviorally equivalent to the multi-selves model. First, we will assume that preferences over decision paths *are* dynamically consistent (although there is no special reason to assume that). This eliminates the possibility of *naivety*, because a consumer who chooses according to a stable preference relation over decision paths behaves as if he knows in period 1 that his second-period choice will maximize $U(A, \cdot)$.

More importantly, the multi-selves and self-control models induce different preferences over menus in period 1. The self-control model allows the consumer to prefer commitment (i.e. a menu with fewer options) even if his choice behavior does not change as a result the commitment. For example, he may prefer the menu $\{z\}$ to the menu $\{z, y\}$ even if he chooses z from the larger menu. This will happen if $v^*(y) > v^*(z)$ and $u^*(z) + v^*(z) > u^*(y) + v^*(y)$. The reason the smaller menu is preferred is that this saves the cost of self-control, which is not required at the singleton menu but is required at the larger menu. In contrast, a multi-selves consumer prefers a smaller menu only if his choice behavior is affected when we remove the commitment.

For instance, let $Z = \{b, i\}$ - where b and i stand for broccoli and icecream. Suppose that $u^*(b) > u^*(i)$ and $v^*(i) > v^*(b)$. This captures a situation in which the consumer is on a diet, and he finds broccoli better than icecream on dietary grounds, yet he finds icecream more tempting than broccoli. It is easy to check that if $u^*(b) - u^*(i) > v^*(i) - v^*(b)$, then:

$$U(\{b\}, b) > U(\{b, i\}, b) > U(\{b, i\}, i) = U(\{i\}, i)$$

such that the consumer chooses b from the menu $\{b, i\}$. The interpretation of these rankings is as follows. When $u^*(b) - u^*(i) > v^*(i) - v^*(b)$, the temptation of icecream is insufficiently strong to override the diet consideration, and therefore the consumer chooses broccoli. However, resisting the temptation to eat icecream requires costly self-control, and for this reason the consumer would like to commit to broccoli rather than having a free choice that demands self-control. That is, self-control is a costly substitute for a perfect commitment device.

Conversely, if $u^*(b) - u^*(i) < v^*(i) - v^*(b)$, then:

$$U(\{b\}, b) > U(\{b, i\}, i) = U(\{i\}, i) > (\{b, i\}, b)$$

such that the consumer chooses i from the menu $\{b, i\}$. The interpretation of these rankings is as follows When $u^*(b) - u^*(i) < v^*(i) - v^*(b)$, icecream is too tempting for the diet to be sustainable. This is a conventional reason for his taste committing to the set $\{b\}$. (Note that the consumer is indifferent between breaking his diet (the outcome $(\{b, i\}, i)$) and not being on a diet in the first place (the outcome $(\{i\}, i)$).)

In general, if $|v^*(z) - v^*(y)| > |u^*(z) - u^*(y)|$ for all distinct $z, y \in Z$, then in period 2 the consumer chooses according to v^*. Consumer behavior in this case is indistinguishable from the multi-selves model in which u^* and v^* represent the first- and second-period selves' preferences. Thus, the multi-selves model is a special case of our model of self-control preferences.

3.1 Implications for Monopolistic Pricing

To see the implications of self-control for pricing behavior, relative to the multi-selves model, let us consider a special setting, in which a monopolist offers one unit of a product. The monopolist is restricted to a simple pricing scheme: it can only charge an entry fee from the consumer in period 1. An alternative is a pair (q, p) where $q \in \{0, 1\}$ is the consumer's consumption decision and p is the payment he ends up making. Assume that both u^* and v^* are quasi-linear in money. Thus, there exist u and v such that $u^*(q, p) = u(q) - p$ and $v^*(q, p) = v(q) - p$. If the consumer chooses not to pay in period 1, then $q = p = 0$. Assume $u(0) = v(0) = 0$ and $u(1) > \max\{0, v(1)\}$.

This example fits the health club example discussed earlier in this chapter, where $q = 1$ (0) represents doing (not doing) physical exercise. Ex-ante, the consumer wants to do physical exercise and he would be willing to pay a positive amount for a mechanism that would force him to work out. However, in period 2 he is tempted not to do any physical exercise. Correspondingly, his willingness to pay for it goes down - possibly below zero.

If the firm could force the consumer to choose $q = 1$ after he pays the entry fee, it could charge $A = u(1)$ for such a commitment device - independently of whether the consumer follows the self-control or multi-selves models. Now suppose that the monopolist cannot monitor whether the consumer chooses $q = 0$ or $q = 1$ after paying the entry fee A. Having paid the fee, the consumer will choose $q = 1$ in period 2 as long as $u(1) + v(1) \geq u(0) + v(0)$. However, when this condition holds, the consumer agrees to pay the entry fee in period 1 only if $u(1) - A + v(1) - A - \max\{v(1) - A, v(0) - A\} \geq 0$. Thus, if $v(1) < 0$, his first-period willingness to pay for entry will be strictly below $u(1)$.

Thus, when the monopolist cannot monitor the consumer's second-period action, it is possible that the consumer will enter and consume the product, but at a lower price than what it could charge for a perfect commitment device. This effect would be impossible under the multi-selves model: if $q = 0$ is a more tempting alternative in period 2 than $q = 1$, then in the absence of a commitment device the consumer will simply refuse to pay any entry fee, anticipating that he will choose $q = 0$ in period 2.

It is crucial for this distinction between the self-control and multi-selves model that the entry fee is sunk. It is were fully refundable - i.e., if the alternative $q = p = 0$ remained feasible in period 2 - the qualitative distinction between the two models would disappear. The reason is that under the multi-selves model (with u and v representing first- and second-period willingness to pay), feasibility of $(0, 0)$ in period 2 means that the firm may find it optimal to lower its entry fee A in order to ensure that $v(1) - A \geq v(0) - 0$. (Of course, this move is profitable only if $v(1) > 0$, whereas in the self-control model, the result that entry fees are lower in the absence of full commitment emerges when $v(1) < 0$.)

Thus, when the entry fee is not sunk, the multi-selves and self-control models generate the same qualitative effect: failure to offer a perfect commitment device leads to a lower entry fee (albeit under conflicting temptation preferences). However, the logic behind the effect is different in the two models. In the multi-selves model, the lower price is designed to maintain $q = 1$ as a second-period incentive compatible action. In the self-control model, it is designed to satisfy the participation constraint. At any rate, to an outside observer the effect is qualitatively the same.

4 A Summary Exercise: Do Temptations Exert an Anti-Competitive Force?

I conclude this chapter with an exercise that applies the self-control model to a simple competitive market setting, in order to demonstrate that the presence of temptations can obstruct competition. The basic idea is that when firms face a population of consumers with diverse tastes, they can act competitively by offering a large variety at attractive prices. However, when some consumers have self-control problems, they may actually dislike variety, and this can weaken competitive forces. As we shall see, the analysis will not depend on whether consumers follow the multi-selves model or the more general self-control model. Nevertheless, the exercise will be presented in terms of the former because it is a generalization of the latter.

Consider the following situation. Two food stores can provide two products, broccoli and icecream (denoted b and i) at zero cost. The stores compete for a measure one of consumers, by choosing simultaneously which products to offer and at what price. Consumers move after the firms make their decisions. They go through a two-stage decision process. First, given the menu of product-price pairs offered by each store, they choose which store to enter. They also have the option of entering none of the stores. Let p_b and p_i denote prices of broccoli and icecream.

Consumers are divided into two groups with different preferences over decision paths:

Group I: $u^*(b, p_b) = 1 - p_b$ and $u^*(i, p_i) = -p_i$
Group II: $u^*(b, p_b) = -p_b$ and $u^*(i, p_i) = 1 - p_i$

As to the consumers' temptation utility v^*, consider first a benchmark situation in which v^* is a constant function for all consumers. $(b, p_b) - v(i, p_i) = 0$ for all consumers. In other words, there are no self-control costs, hence consumer behavior collapses to a standard model of rational choice with heterogenous preferences. For all consumers, the u^* and v^* values associated with not entering any store is 0.

In this environment, there is a unique Nash equilibrium, in which both stores offer both products at zero price. This is a perfectly competitive outcome. The logic behind this result is standard: as long as one store offers any product at a strictly positive price, the other firm has an incentive to offer the same product at a slightly lower price and win the entire group of consumers who prefer this product.

Now consider an alternative scenario, which differs from the benchmark scenario only in the v^*-values of the first group of consumers, so that $v^*(i, p_i) - v^*(b, p_b) = \delta > 1$. The interpretation is that consumers whose commitment utility ranks broccoli above icecream are simply on a healthy diet and they find icecream more tempting than broccoli. Note that temptation utility ignores product prices in this example.

We will now show that in this alternative environment, there is a Nash equilibrium in which store 1 offers only b at a price $p_b = 1$, while store 2 offers only i at a price $p_i = 1$. First, note that under this strategy profile, group I is indifferent between shopping at store 1 and not shopping at all, while group II is indifferent between shopping at store 2 and not shopping at all. Since the two consumer groups are equal in size, no store has an incentive to *replace* the product they offer with the other product and undercut the price that the competing store charges for the latter. Therefore, the only potentially profitable deviation is by *adding* a product to the store' menu. Let us show that such a deviation is unprofitable:

- Suppose that store 1 adds i to its menu, at a price p_i. It may also change the price it charges for broccoli p_b. In order for the deviation to be profitable, p_b cannot be negative. In order to attract consumers away from store 2, it must be the case that $p_i < 1$. However, this makes store 1 unattractive for group-I consumers. If they enter the store and buy broccoli, their payoff is $1 - p_b - \delta < 0$. If they enter the store and buy icecream, their payoff is $-p_i < 0$. Thus, group-I consumers prefer not to enter store 1 as a result of this deviation. It follows that the deviation cannot be profitable, because store 1's payoff is $\frac{1}{2}p_i$, whereas prior to the deviation it was $\frac{1}{2}$.

- Suppose that store 2 adds b to its menu, at a price p_b. It may also change the price it charges for icecream p_i. In order for the deviation to be profitable, p_b

cannot be negative. However, as long as icecream is on store 2's menu, group-I consumers prefer not to enter it. If they enter the store and buy icecream, their payoff is $-p_i$. If they enter and buy broccoli, their payoff is $1 - p_b - \delta < 0$. Thus, store 2 cannot attract group-I consumers without lowering its profits.

Thus, in this Nash equilibrium stores specialize in different products and offer them at monopolistic prices. The reason for this equilibrium differentiation is that in order to engage in profitable price competition, firms must offer both products. However, such variety is anathema to consumers on a diet, and therefore competitive forces are weakened. If each store could set up a second branch, so that each branch would specialize in a different product, then the ordinary competitive force would be restored because firms would be able to offer broccoli to dieting consumers without requiring them to exert any self-control cost.

As mentioned above, the equilibrium specialization result does not depend on whether not $|u^*(i, p_i) - u^*(b, p_b)| \geq |v^*(i, p_i) - v^*(b, p_b)|$ - i.e., on whether consumers obey the multi-selves model or the more general self-control model. Finding general conditions for the two models to yield identical qualitative predictions in I.O. models is an interesting open problem.

5 Summary

In this chapter we examined considerations that arise when firms have limited ability to monitor the consumer's second-period action. Imperfect monitoring limits both the commitment devices that firms can offer to sophisticates and the exploitative, betting contracts that they can offer to naifs.

- When firms are restricted to two-part tariffs, we should expect to observe deviations from marginal-cost pricing. The direction of the deviation depends on whether second-period preferences display greater or lower willingness to pay than first-period preferences, but not on whether consumers are naive or sophisticated.

- Under the standard conditions on preferences that would justify two-part tariffs as being optimal in a standard pricing model, they are optimal the consumer is sophisticated, but not when he is naive. Moreover, the naive consumer is better off (in terms of his first-period preferences) if firms face this restriction to two-part tariffs.

- Renegotiation between the firm and the consumer in period 2, or market competition with other suppliers in period 2, restrict considerably the commitment devices that firms can offer consumers. However, they do not upset the betting contracts signed with naifs.

- When firms have limited ability to offer perfect commitment devices, self-control can serve as a substitute mechanism. However, since this mechanism may be costly, it will lower the prices that firms can charge from sophisticated consumers.

- When a population of consumers contains a group of consumers with self-control costs, this can make the market outcome less competitive than in the absence of such considerations. However, this conclusion is unique to the self-control model, and can also exist when consumers have no self-control and follow the multi-selves model. In general, although the two models are distinguishable at the individual behavior level, they are hard to distinguish in terms of the market equilibrium outcomes they give rise to

6 Bibliographic Notes

The characterization of optimal two-part tariffs is based on DellaVigna and Malmendier (2004). The discussion of renegotiation and spot market competition builds on Koszegi (2005). The self-control model is due to Gul and Pesendorfer (2001). More general formulations were proposed by Dekel, Lipman and Rustichini (2006). The two market applications of the self-control model are free variations on Esteban and Miyagawa (2006,2007).

References

[1] Dekel, E., B. Lipman and A. Rustichini (2006), "Temptation-Driven Preferences", mimeo, Boston University.

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

[3] Esteban, S. and E. Miyagawa (2006), "Temptation, Self-Control, and Competitive Nonlinear Pricing", *Economic Letters* **90**, 348-355.

[4] Esteban, S. and E. Miyagawa (2007), "Optimal Menus of Menus with Self-Control Preferences", mimeo.

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

[6] Koszegi, B. (2005), "On the Feasibility of Market Solutions to Self-Control Problems", *Swedish Economic Policy Review* **12**, 71-94.