

NBER WORKING PAPER SERIES

OVERESTIMATING SELF-CONTROL:
EVIDENCE FROM THE HEALTH CLUB INDUSTRY

Stefano DellaVigna
Ulrike Malmendier

Working Paper 10819
<http://www.nber.org/papers/w10819>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2004

An earlier version of this paper was distributed under the title “Self-Control in the Market: Evidence from the Health Club Industry.” We are particularly grateful to Edward Glaeser, Lawrence Katz, and David Laibson for their invaluable support. We thank the editor, an anonymous referee, Marios Angeletos, George Baker, Rachel Croson, Rajeev Dehejia, Oliver Hart, Caroline Hoxby, Daniele Paserman, Antonio Rangel, Andrei Shleifer, Jeremy Tobacman, Klaus Wertenbroch, Justin Wolfers, and the participants of seminars in CREST, Dartmouth College, Harvard, Michigan (Ann Arbor), Northwestern, NYU, UC Berkeley, UCLA, at HBS, Kellogg, Haas, at the Ente Einaudi, at the SITE 2001 and 2002, and at the Eastern Economic Association 2002 for their comments. We thank Bryan S. Graham for providing Mathematica code and Tobias Adrian, Augustin Landier, and Sendhil Mullainathan for inspiring conversations at the beginning of this project. Tricia Glynn, Camelia Kuhnen, Scott McLinn, Boris Nenchev, and Nikita Piankov provided excellent research assistance. For financial support, DellaVigna thanks Bank of Italy, Malmendier thanks the DAAD. The views expressed herein are those of the author(s) and not necessarily those of the National Bureau of Economic Research.

©2004 by Stefano DellaVigna and Ulrike Malmendier. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Overestimating Self-Control: Evidence from the Health Club Industry
Stefano DellaVigna and Ulrike Malmendier
NBER Working Paper No. 10819
September 2004
JEL No. D12, D18, D21, D23, D91, D92, L11, L83

ABSTRACT

Experimental evidence suggests that people make time-inconsistent choices and display overconfidence about positive personal attributes. Do these features affect consumer behavior in the market? To address this question we use a new panel data set from three US health clubs with information on the contract choices and the day-to-day attendance decisions of 7,978 health club members over three years. Members who choose a contract with a flat monthly fee of over \$70 attend on average 4.8 times per month. They pay a price per expected visit of more than \$17, even though a \$10-per-visit fee is also available. On average, these users forgo savings of \$700 during their membership. We review many aspects of the consumer behavior, including the interval between last attendance and contract termination, the survival probability, and the correlation between different consumption choices. The empirical results are difficult to reconcile with the standard assumption of time-consistent preferences and rational expectations. A model of time-inconsistent agents with overconfidence about future patience explains the findings. The agents overestimate the future attendance and delay contract cancellation whenever renewal is automatic. Salesman pressure and overestimation of future efficiency are the leading alternative explanations.

Stefano DellaVigna
University of California, Berkeley
Department of Economics
549 Evans Hall, #3880
Berkeley, CA 94720-3880
sdellavi@econ.berkeley.edu

Ulrike Malmendier
Graduate School of Business
518 Memorial Way
Stanford University
Stanford, CA 94305-5015
and NBER
ulrikem@stanford.edu

^{fs}
“Saturday 31 December. New Year’s Resolutions. I WILL [...] go to the gym three times a week not merely to buy sandwich.” (Fielding, 1999. *Bridget Jones’ Diary: A Novel*)

A few months later: “Monday 28 April. [...] Gym visits 0, no. of gym visits so far this year 1, cost of gym membership per year £370; cost of single gym visit £123 (v. bad economy).” (Fielding, 2001. *Bridget Jones: The Edge of Reason*)

1 Introduction

Experimental evidence calls into question the assumptions of time consistency (Kirby and Herrnstein, 1995; Loewenstein and Prelec, 1992; Thaler, 1981) and rational expectations about own abilities (Larwood and Whittaker, 1977; Svenson, 1981). Given the importance of the economic implications, it is crucial to know whether time inconsistency and overconfidence, as displayed in laboratory settings, affect consumer behavior in the market. Recent empirical studies point to market evidence of time inconsistency in the fields of consumption (Angeletos et al., 2001), addiction (Gruber and Mullainathan, 2002), and job search (DellaVigna and Paserman, 2000). Market evidence of overconfidence appears in takeovers and corporate investment (Malmendier and Tate, 2002 and 2003).¹

This paper contributes to the empirical evidence on time inconsistency and overconfidence by analyzing a simple yet economically significant decision, enrollment and attendance in a health club. The simplicity and familiarity of the contractual choice allow us to address alternative interpretations of the previous literature, such as limited cognitive abilities. We employ a new panel data set from three US health clubs with 7,978 members over three years. Unlike most data sets on consumption behavior, this data set documents both the purchase of a commodity—enrollment in the health club and membership renewal—and the actual consumption—health club attendance. We complement this data set with a survey of health club members, a survey of health club companies, and a field experiment.

A crucial feature of this data set is the presence of a menu of contractual options. Consumers can choose between two flat-rate contracts—a monthly contract and an annual contract—and a pay-per-visit option for \$10. The monthly contract is automatically renewed from month to month until the consumer cancels. The annual contract, instead, expires after twelve months and the consumer has to explicitly renew it. We exploit the variation in the per-usage pricing and in the renewal procedures to identify features of consumer preferences and beliefs.

We establish nine stylized facts, summarized in Table 1. We first take advantage of the presence of both flat-rate and pay-per-visit contracts in the menu of memberships. We show that health club members who choose a monthly contract with a flat fee of over \$70 attend on

¹Aiely and Wertenbroch (2002), Fang and Silverman (2001), and Gruber and Koszegi (2001) also present field evidence on time inconsistency.

average 4.8 times per month in the first six months. They pay a price per expected visit in excess of \$17, even though a \$10-per-visit fee is also available. Only 20 percent of the users in the monthly contract pay less than \$10 per visit ex post (Stylized Fact 1). Comparable findings hold for the first year of the annual contract. Using a survey, we document that health club users overestimate their future usage by more than 100% percent (Stylized Fact 2).

The next stylized facts relate to the contractual choice over time. Users with a monthly contract accumulate an average gap of 2.29 full consecutive months between the last attendance and contract termination, at a monetary cost of \$185 (Stylized Fact 3). We compare the renewal behavior under the monthly and annual contract. Under the monthly contract the initial attendance is 9.5 percent lower (Stylized Fact 4), but the likelihood of still being enrolled after one year is 12.5 percent higher (Stylized Fact 5). Surprisingly, the contract that allows *more* freedom to cancel has a *higher* share of agents renewing the contract, even though long-term enrollment is cheaper under the annual contract. This difference is particularly high for agents with low attendance in the first 12 months (Stylized Fact 6).

We also consider the dynamics of average attendance for agents initially enrolled in the annual contract. Average monthly attendance in the first year is 46 percent lower than the attendance of those who re-enroll in the second year (Stylized Fact 8). The pattern is reversed for users initially enrolled in the monthly contract. Average monthly attendance in the first six months is significantly higher than in any of the subsequent six-month periods (Stylized Fact 7). Finally, users who pay a high price per attendance in the monthly contract display a longer gap between last attendance and contract termination (Stylized Fact 9).

Stylized Facts 1 through 9 are hard to reconcile with a standard model of time-consistent decision-makers with rational expectations. In Sections 3.3 and 4.3 we discuss a number of potential explanations of these facts, including high transaction costs of payment per usage, risk aversion, overestimation of net benefits of attendance, and salesman techniques. While each of these interpretations explains some of the stylized facts, none provides a unifying explanation for the overall pattern of the evidence.

In Section 3.1 we present a model of partially naive time-inconsistent agents (Strotz, 1956; Phelps and Pollak, 1968; Laibson, 1997; O'Donoghue and Rabin, 1999 and 2001) that organizes all empirical findings. Agents with these features pay more than \$10 per expected visit (Stylized Fact 1) for two reasons: they purchase a commitment device that increases future attendance and they overpay due to overestimation of future attendance. Naive, time-inconsistent agents also display a status quo effect (Samuelson and Zeckhauser, 1988; Madrian and Shea, 2001) and delay membership cancellation under the monthly contract (Stylized Fact 3). Since they expect to be more patient in the near future, they delegate cancellation to later selves in the (irrational) belief that these later selves will cancel (Akerlof, 1991; O'Donoghue and Rabin, 2001). The delay of cancellation does not occur under the annual contract, since the annual contract expires automatically after 12 month. As a result, the share of agents enrolled after

12 months is higher under the monthly contract than under the annual contract (Stylized Facts 5 and 6). Under the annual contract, selective exit of users with low attendance explains the higher average attendance for stayers than for the initial group of members (Stylized Fact 7). Under the monthly contract, delay of contract cancellation explains the decrease in average attendance, as negative shocks accumulate (Stylized Fact 8). Finally, in the presence of heterogeneity, naive individuals who pay a higher price per attendance also have a longer cancellation gap (Stylized Fact 9).

We highlight two alternative explanations that capture most of the facts. First, consumers may underestimate both the net costs of attendance and the cancellation costs. These consumers are overconfident about their future attendance and display a status-quo bias. This model explains the stylized facts for the same reasons that the model of naive time-inconsistent agents does. A difference is that this model predicts no demand for commitment devices. A second explanation is salesman pressure. Salespeople at the health club may use social pressure to induce agents to first purchase flat-rate contracts, and then remain enrolled even for low levels of attendance. This explanation, however, does not address the survey evidence on overestimation of future attendance (Stylized Fact 2).

This paper provides market evidence of systematic and large deviations of consumer behavior from the standard model. These deviations occur in a familiar market setting, the choice between flat-rate and payment per usage. In the health clubs of our sample, the average non-subsidized user chooses the monthly contract, and by doing so forgoes savings of about \$700 per membership, out of a total amount of about \$1,500 paid to the health club. The results of this study are likely to generalize to the 32.8m Americans who exercise in one of the 16,983 US health clubs. Therefore, both in terms of monetary magnitude and in terms of population involved, the non-standard behavior has a significant economic impact in the health club industry. Our interpretation of this deviation relies on time inconsistency and naiveté.

While the systematic deviations from the standard model are the core result of the paper, we would like to emphasize two other themes. First, despite deviations from the standard models, health club users are responsive to standard economic forces. We find strong evidence that individuals learn over time, and that they switch toward the contract that is more appropriate given their attendance. Average attendance among stayers in the annual contract is substantially higher than in the first year (Stylized Fact 7). In addition, the observed initial sorting between the monthly and annual contract conforms to the standard predictions for agents with heterogeneous attendance costs (Stylized Fact 4).

Second, while the paper focuses on the consumer side of the market, it has implications for the industrial organization of flat-rate pricing² and of automatic renewal. We show that,

²The previous studies on flat-rate and pay-per-visit pricing had focused almost exclusively on the telecommunication industry (Miravete, forthcoming).

in contracts with automatic renewal, a small cancellation cost induces substantially higher survival rates. Arguably, this finding can explain the frequency of contracts with automatic renewal in other industries such as the newspaper, credit card, and mail order industry. In DellaVigna and Malmendier (2002) we explore the general implications of time inconsistency and naïveté for firm pricing.

The remainder of the paper is organized as follows. In Section 2 we introduce the main features of the health club data set. In Section 3 we consider the contractual choice at enrollment and in Section 4 we analyze the contractual choice over time. In both Sections, a simple model introduces the Stylized Facts and the discussion of the interpretations. In Section 5 we evaluate the size of the deviations from the standard model. Section 6 concludes.

2 Health club data set

Health club industry. As of January 2001, 16,983 clubs were operating in the US. The industry revenues for the year 2000 totalled \$11.6bn. The memberships in the same period summed to 32.8m, up from 17.4m in 1987. Fifty-one percent of the users were members in commercial health clubs, while thirty-four percent were members in non-profit facilities. Only the market leader Bally Total Fitness with \$1,007m revenues and 4m members is publicly traded. Few companies operate in more than 10 states. Ownership concentration is in the 10th percentile of US industries.

Data set. We collected a new panel data set from three health clubs located in New England, which we label clubs 1, 2, and 3. The data set contains information on the contractual choices and the day-to-day attendance of users that enrolled after April 1, 1997. The sample period lasts until August 24, 2000 for club 1 and until March 13, 2001 for clubs 2 and 3. The day-to-day record of usage is made available by the technology regulating the access to these health clubs, described below. The panel of contractual choices comes from the billing records. Each entry in the accounting data specifies the price paid for the transaction and a 4-letter code. This code allows us to track the membership type—standard, student, family, corporate—as well as details like the subsidizing company (if any).

Several companies located near the clubs subsidize their employees’ attendance. For these corporate members, the health club receives part of the membership payments directly from the firms, with the remainder being paid by the members. The health club informs the companies periodically about the number of employees enrolled and their attendance. This creates incentives for the health club to record attendances accurately or, possibly, to overrecord them.

Contractual menu. We conducted a survey of the 100 health clubs in the metropolitan area of Boston to document the contract design in the industry.³ Health clubs offer three

³For details on the survey, see DellaVigna and Malmendier (2002).

options to attend. 85 clubs offer a monthly contract. A monthly fee is automatically debited each month to a credit card or bank account until the user cancels the membership. 90 clubs offer an annual contract that involves the payment of an annual fee. Both monthly and annual contracts have an initiation fee and no fee per visit. Finally, 82 clubs offer a pay-per-visit option, often in the form of a 10-visit pass. Health clubs 1 and 2 in our sample offer the three types of contract with the following additional features.⁴

1. The monthly contract has a monthly fee ranging between \$85 (standard level) and \$85 (discounted level). Non-corporate users also pay an initiation fee ranging from \$0 (in promotional periods) to \$150. Corporate users generally pay an out-of-pocket monthly fee between \$19 and \$65, as a function of the subsidy paid by their company, and no initiation fee. Cancellation can be done in person at the club or by sending a written note.⁵ If cancellation takes place before the 10th of the month, no further fees are due, and the users can attend until the end of the month. Members who cancel after the 10th have to pay the fee for the next month and can attend until the end of the following month.
2. The annual contract charges up-front 10 times the applicable monthly fee, for a standard fee of \$850. Users thus get a discount of 2 months out of 12 in exchange for a yearly commitment. The initiation fee is the same as under the corresponding monthly contract. At the end of the year, the contract expires and members who wish to stay enrolled have to sign up again, either for an annual or for a monthly contract. In order to encourage renewal, the club sends out a reminder card one month before the contract expires.
3. The pay-per-visit system offers two options, either to pay \$12 per visit or to purchase a 10-visit pass for \$100. Transaction costs for the 10-visit pass are small. Users provide basic demographic information and receive a card for ten visits. Unfortunately, attendance is not tracked for the pay-per-visit users.

Users of club 3 face the same menu of contracts with lower prices and slightly different services. The monthly fee ranges from \$13 to \$52, and the initiation fee is at most \$50. The annual fee in the annual contract equals 10 times the corresponding monthly fee. The pay-per-visit options are a \$10 fee per visit, and a \$80 pass for 10 visits. Finally, the enrollment in a monthly or an annual contract does not include the provision of towels.

Under these three membership types, users deposit their cards in a basket at the front desk when they enter. While they are exercising, a health club employee swipes them (marks the

⁴Contracts for one to six months with automatic expiration are also available. We do not include them in our analysis, since they are typically targeted towards occasional summer users. We also remove from the sample free limited-time memberships that are occasionally given to employees of the subsidizing companies.

⁵Some users cancel by discontinuing the payments to the health club.

visit for the 10-visit passes), and users pick them up when they exit. This method guarantees a high recording precision even during peak hours. The three contracts give right to the same services, i.e., a temporary locker, towels, and access to the equipment. Also, both the monthly and the annual contract allow members to “freeze” (suspend) their membership for three months per year⁶. Users with a monthly contract do not have to pay their monthly fee during the freezing period. Annual members get additional usage time after the twelve months of the original membership.

The cancellation policy of the monthly contract involves a monthly deadline for cancellation (the 10th of the month). To find out whether consumers are aware of the exact deadline, we conducted a survey among health club users. Since it was not feasible to survey members of the three clubs in our sample, we collected data from random consumers in a mall in California. 48 consumers who stated that they attended a health club completed a survey⁷. Out of these, 32 individuals indicated a monthly membership and were then asked by which day they have to cancel ‘in order to avoid paying the next monthly fee.’ All but one stated that they did not know the answer or reported “cancellation any time (30 days in advance).” A follow-up with the health clubs revealed that for 20 out of 27 individuals (whom we could match to a health club), the clubs indeed had a monthly deadline. Therefore, out of 20 individuals enrolled in 15 different clubs with monthly deadlines, only one knew the relevant date. Building on this evidence, we will assume that the typical individual is uncertain about the monthly deadline for cancellation (see Section 4.1).

Sample construction. We match the information on attendance and on contract choice in the three clubs to form a longitudinal data set with monthly observations, covering the period from April 1997 to August 2000 (club 1) and to March 2001 (clubs 2 and 3). Our analysis focuses on *enrollment spells*. A spell starts whenever an individual enrolls (or re-enrolls) in the club and ends whenever the individual quits. We define spells to be censored if either the enrollment is ongoing at the end of the sample period, or the individual switches to a short-term contract or receives a promotional membership. Individuals have multiple spells if they quit the club and re-enroll at some later date.

The initial sample includes 11,605 individuals. We drop individuals who were never enrolled in either a monthly or an annual contract (2,978 individuals). We eliminate spells with serious inconsistency in the billing data (132 spells). We also exclude users with a family membership to avoid issues regarding the joint consumption of the services (295 spells). Finally, in order to limit the sample to first-time users of these clubs, we drop users who had a free or a seasonal membership before they chose a monthly or an annual contract (293 individuals).

Enrollment spells. This leaves us with a sample of 7,978 individuals and 8,615 enrollment spells. In the paper, we use the sample ‘First spell’, which includes only the first enrollment

⁶Monthly users can also quit for up to three additional months without repaying the initiation fee.

⁷The Survey protocol is available in the working paper version in DellaVigna and Malmendier (2003).

spell for each individual. As Row 1 of Table 2 shows, Club 1 has 19 percent more members than club 2, and more than twice as many members as club 3. The percentage of completed spells (above 60 percent) is similar across the clubs. Of the 7,978 individuals enrolled in any club, 87 percent choose a monthly membership as their first contract. Health club members rarely change the type of contract they initially enroll in. *We* also use the sample ‘First spell and no subsidy’, a restriction of the sample ‘First spell’ to unsubsidized memberships. We consider a membership to be unsubsidized if, over the whole spell, the average out-of-pocket fee exceeds \$70 per month for enrollment in a monthly membership and \$700 per year (\$58 per month) for enrollment in an annual membership. This smaller sample includes 1,120 spells (14.03 percent of the full sample).

Descriptive statistics. In clubs 1 and 2, the average amount spent per spell is about \$550, and the average fee per month ranges between \$43 and \$53. For corporate users, these are the out-of-pocket payments and do not include the subsidies paid by the sponsoring firms. In club 3 the amounts are substantially lower, since the contracts are cheaper. In the sample ‘First spell and no subsidy’ (Columns 7 and 8), these amounts are 20 to 60 percent higher. The initiation fee averages \$4 in the sample ‘First spell’ since 86 percent of users do not pay it. In the ‘First spell’ sample, individuals with a monthly contract attend on average 4 times per month, and individuals with an annual contract attend on average 4.3 times per month. Attendance in club 1 (Column 1) is somewhat higher than in the other clubs. Freezing of a contract is rare in all the clubs. The bottom part of Table 2 displays the available demographic controls. Users are somewhat more likely to be male than female and are on average in the early thirties. Corporate memberships account for 50 percent of the sample, while student memberships account for only 2 percent. Additional information on the data set construction is available in the Data Appendix.

3 Contract choice at enrollment

3.1 Model

In this Section we provide a simple model of contract choice for consumers attending a health club. The contractual options mirror those offered in the three clubs in the sample. The model captures the standard case of time-consistent decision makers with rational expectations. The framework also embeds alternative preference and belief specifications, such as time inconsistency, overestimation of future benefits, and high transaction costs. In this Section and in Section 4.1 we highlight eight testable predictions for the standard model. For a more formal presentation, see DellaVigna and Malmendier (2003).

Contractual menu. A contract specifies a sign-up fee L' , a price per usage p' , and the number of periods (days) T' of contract duration. The firm offers consumers a choice between

flat-fee contracts, characterized by $p' = 0$, and pay-per-visit contracts, characterized by $L' = 0$. The fee p' in pay-per-visit contracts incorporates the transaction costs. Flat-fee contracts are either *monthly* ($T' = 30$ days) or *annual* ($T' = 360$ days). The timing is as follows. At time 0, the consumers choose between a flat-rate (monthly or annual) and a pay-per-visit contract. They pay the lump-sum fee L' at $t = 1$ and have the option, at $t = 1, \dots, T'$, to attend the health club or to pursue their best alternative activity. If they attend at time t , they pay the fee p' at time t . At $t = T'$, the agents choose again between the different contracts, with the lump-sum fee L' due at $T' + 1$, and so on.

Payoffs of attendance. If the agents attend the health club at time t they incur an immediate cost c at time t and reap benefits $b > 0$ at time $t + 1$. If they do not attend, they attain the payoff from the best alternative activity, which we normalize to 0. We interpret c as the effort cost and b as the (expected) net present value of all the future benefits from better health and improved fitness. We assume that the cost of attendance c is uncertain at sign-up, with a distribution G . The value of c is realized after sign-up and is constant thereafter.

Intertemporal preferences. The benchmark assumption is that agents have time-consistent preferences with daily discount factor δ . We embed this assumption in a model that allows for quasi-hyperbolic preferences (Phelps and Pollak 1968, Laibson 1997, O'Donoghue and Rabin 1999). The discount function for time s , evaluated at period t , equals 1 for $s = t$ and $\beta\delta^{s-t}$ for $s = t + 1, t + 2, \dots$, with $\beta = 1$ in the baseline case and $\beta < 1$ in the case of time inconsistency. The present value of a flow of future utilities $(u_s)_{s \geq t}$ as of time t is

$$u_t + \beta \sum_{s=t+1}^{\infty} \delta^{s-t} u_s. \quad (1)$$

We can interpret β as the parameter of short-run discounting and δ as the parameter of long-run discounting. If β is smaller than 1, the individuals exhibit time-varying discounting. The discount factor between the present period and the next period is $\beta\delta$, while the discount factor between any two periods in the future is simply δ . The difference between the short-run and the long-run discount factors generates *time inconsistency*.

We also consider the case of consumers who overestimate their time consistency. *Partially naive time-inconsistent* agents with parameters $(\beta, \hat{\beta}, \delta)$ (O'Donoghue and Rabin, 2001) expect (erroneously) to have the discount function $1, \hat{\beta}\delta, \hat{\beta}\delta^2, \dots$ with $\beta < \hat{\beta} \leq 1$ in all future periods. The individuals thus anticipate that they will have quasi-hyperbolic preferences in the future, but they overestimate the future parameter of short-run discounting. The difference between the perceived and actual future short-run discount factor $\hat{\beta} - \beta$ reflects the *overconfidence* (naiveté) about future self-control. If the consumers have rational expectations, $\hat{\beta}$ equals β .

Attendance decision. At enrollment ($t = 0$), the agents assign discounted net utility $\beta\delta^t(\delta b - p - c)$ to attendance at time t , and utility 0 to the best alternative activity. Thus, they would like to attend in a future period t , upon learning c , if $c \leq \delta b - p$.

While *time-consistent rational* agents attend as often as they wish, *time-inconsistent* agents attend less often. At the moment of deciding whether to attend, the discounted payoff of attendance is $\beta\delta b - p - c$. Therefore, at each t they attend only if $c \leq \beta\delta b - p$, i.e., with probability $G(\beta\delta b - p)$. The smaller is β , the larger is the difference between desired and actual attendance probability $G(\delta b - p) - G(\beta\delta b - p)$ and the more serious is the time inconsistency.

Partially naive time-inconsistent individuals with $\beta < \hat{\beta}$ are not fully aware of their time inconsistency. At time 0 they overestimate the probability with which they will attend in future periods. They expect to attend if $c \leq \hat{\beta}\delta b - p$, i.e., with probability $G(\hat{\beta}\delta b - p)$. The larger is the difference $\hat{\beta} - \beta$, the larger is the overestimation of the future attendance probability $G(\hat{\beta}\delta b - p) - G(\beta\delta b - p) \geq 0$.

Contractual choice. At time 0, consumers who sign a contract (T', L', p') expect to attain the net benefit

$$\beta\delta \left[-L' + \frac{1 - \delta^{T'}}{1 - \delta} \int_{-\infty}^{\hat{\beta}\delta b - p'} (\delta b - p' - c) dG(c) \right]. \quad (2)$$

The agents pay a fee L' , independently of attendance, and attain utility $\beta\delta^t(\delta b - p' - c)$ from attending at time t , with $1 \leq t \leq T'$ (notice the $\hat{\beta}$ in the integral). Using expression (2) it is easy to show that at $t = 0$ agents prefer a flat-rate contract $(T', L, 0)$ to a pay-per-visit contract $(T', 0, p)$ if the flat fee L is smaller than the willingness to pay for the flat-fee contract:

$$L \leq \frac{1 - \delta^{T'}}{1 - \delta} \left[\int_{-\infty}^{\hat{\beta}\delta b - p} p dG(c) + \int_{\hat{\beta}\delta b - p}^{\hat{\beta}\delta b} (\delta b - c) dG(c) + \int_{\hat{\beta}\delta b}^{\infty} 0 dG(c) \right]. \quad (3)$$

Whenever the net value per visit is high ($c \leq \hat{\beta}\delta b - p$), the consumers expect to attend regardless of whether they have to pay a price p (in the pay-per-visit contract) or not (in the flat-rate contract). They are willing to pay $\min(\delta b - c, p) = p$, the price of one visit under the pay-per-visit scheme. Whenever the value of the visits is positive but lower ($\hat{\beta}\delta b - p \leq c < \hat{\beta}\delta b$), the users attend only if $p = 0$, i.e., under the flat-rate contract. In this case they are willing to pay up to the value of a visit under the flat-rate contract, $(\delta b - c)$. Finally, for low value of visits ($c > \hat{\beta}\delta b$), the users attend under neither contract, and the willingness to pay is zero.

Using (3) and inequality $\int_{\hat{\beta}\delta b - p}^{\hat{\beta}\delta b} (\delta b - c) dG(c) \leq [(1 - \hat{\beta})\delta b + p] \cdot [G(\hat{\beta}\delta b) - G(\hat{\beta}\delta b - p)]$, we prove the following Proposition.

Proposition 1. *For agents that prefer a flat-rate contract $(T', L, 0)$ over a pay-per-visit contract $(T', 0, p)$,*

$$\begin{aligned} \frac{(1 - \delta) T'}{1 - \delta^{T'}} L &\leq p \cdot T' G(\beta\delta b) \\ &+ (1 - \hat{\beta})\delta b \cdot T' \left[G(\hat{\beta}\delta b) - G(\hat{\beta}\delta b - p) \right] \\ &+ p \cdot T' \left[G(\hat{\beta}\delta b) - G(\beta\delta b) \right]. \end{aligned} \quad (4)$$

For the baseline case of agents with standard *time-consistent* preferences and rational expectations ($\beta = \hat{\beta} = 1$), Proposition 1 implies the following testable prediction.

Prediction 1. (Price per expected attendance at enrollment) *Agents who choose a flat-rate contract $(T', L, 0)$ over the pay-per-visit contract $(T', 0, p)$ attend frequently enough such that*

$$\frac{(1 - \delta)T'}{1 - \delta^{T'}} L / T' G(\beta\delta b) \leq p \quad (5)$$

Time-consistent agents are willing to pay at most p per expected attendance, the price of the pay-per-visit contract. The adjustment term $(1 - \delta)T' / (1 - \delta^{T'})$ derives from the fact that the fee L is paid up-front at $t = 1$, while the price p is paid between periods 1 and T' .

Time-inconsistent agents may instead pay more than p for two distinct reasons. To the extent that they are *sophisticated* ($\hat{\beta} < 1$), they value the flat-rate contract as a commitment device to overcome the short-run impatience. The second term in (4) is the price that the consumers are willing to pay for this commitment: $(1 - \hat{\beta})\delta b$ is the additional utility of each extra attendance on top of the utility for the future impatient self. This term multiplies the expected increase in attendance $T'[G(\hat{\beta}\delta b) - G(\hat{\beta}\delta b - p)]$ induced by the zero per-visit price. Second, to the extent that the agents are *naïve* about the future time preferences ($\beta < \hat{\beta}$), they overestimate the number of attendances. The additional willingness to pay (relative to time-consistent agents), captured in the third term in (4), equals the overestimation of attendance $T'[G(\hat{\beta}\delta b) - G(\beta\delta b)]$ times the willingness to pay for an attendance, p .

As a result, Prediction 1 does not need to hold for time-inconsistent agents, though it may still hold if time inconsistency is limited (β close to 1). For instance, consumers may overestimate their future attendance, but still attend often enough under the flat-rate contract to pay less than \$10 per visit.

Since condition (5) is only a necessary condition for the standard model to hold, Prediction 1 provides a conservative test of the standard model of time-consistent agents with rational expectations. Moreover, we are neglecting initiation fees and cancellation costs for the flat-rate contracts, which further biases the test in favor of the standard model.

A second, basic test for the standard model is whether consumers have rational expectations about their attendance probability. Denote by \hat{G} the distribution of costs that the agents forecast as of time 0. Rational expectations implies $\hat{G} = G$.

Prediction 2. (Forecast of attendance) *The forecast of attendance $\hat{G}(\hat{\beta}\delta b)$ should equal the average attendance $G(\beta\delta b)$.*

3.2 Empirical analysis

We test Prediction 1 using the sample of users enrolled in an unsubsidized flat-rate membership in clubs 1 and 2. We analyze separately users in club 3 given the lower fee per visit. We consider

the 10-visit pass to be the empirical counterpart of the pay-per-visit contract, so $p = \$10$.⁸

Monthly contract. For users initially enrolled in a monthly contract, we compute the price per expected attendance for each month. We limit the analysis to the first 6 months of tenure to target inexperienced users. We use the sample ‘First spell and no subsidy’ (912 spells) to ensure comparability to standard health clubs with no corporate subsidy.

The first Column in Table 3 reports the average monthly fees in months 1 through 6, with standards errors in parentheses. The sample for month t consists of users who initially enrolled in a monthly contract and have had a continuous history of membership under either a monthly or an annual contract. Consumers who cancel or are censored drop out of the sample. For users who switch to an annual contract, the monthly fee is the monthly share of the annual fee. The average monthly fee exceeds \$80 in all months, except in the joining month which is typically pro-rated, and in month 3, a promotional free month for 18.6 percent of the sample. The average number of visits for users in the t^{th} month of tenure (Column 2) declines from 5.45 in month 2 to 4.32 in month 6 (where, again, month 1 covers only part of a month).

The third Column in Table 3 presents the ratio of the average fee in month t (Column 1) and the average attendance in month t (Column 2). This ratio is the estimated price per expected attendance for month t . In each of the six months we reject the hypothesis that the measure is smaller than \$10. The estimate ranges between \$14 and \$16 in the first three months and is higher than \$17 in the subsequent three months. As a summary measure, we compute the ratio of average monthly payment (Column 1) and average monthly attendance (Column 2) in the first six months across all individuals.⁹ The resulting price per average attendance in the first six months of enrollment equals \$17.13, well above \$10.

In addition to averages, we consider also the distribution of these measures in the first six months (Table 4). We measure the price per attendance as the ratio of total attendance over total payment in the first six months of membership in a monthly contract. In this sample, only 20 percent of the individuals pay less than \$10 per visit. The remaining 80 percent would have saved money choosing the pay-per-visit contract.

Annual contract. We also test Prediction 1 on the sample of users who chose an annual contract at enrollment. We restrict the sample ‘First spell and no subsidy’ to users who joined the club at least 14 month before the end of the sample period. This ensures that we observe the annual contract in its entirety.¹⁰ The final sample consists of 145 spells.

⁸The (hypothetical) average price per average attendance from using the 10-visit pass, given the distribution of attendance for users enrolled with the monthly and the annual contract, is \$10.86. The benefits of a lower price relative to the \$12 per-visit fee outweigh the losses from unused coupons.

⁹For each individual, we compute the average over all available months until the sixth, with the exception of miscoded months and months with freezing. When averaging across individuals, we weigh all individuals equally, independent of tenure.

¹⁰We exclude 3 annual contracts that are terminated before the 12th month. Health clubs are required to accept cancellations for medical reasons or for relocation more than 25 miles away from the clubs.

The bottom row of Table 3 presents the estimation results. The sample average of the monthly share of the annual fee for the first year (Column 1), adjusted for discounting, is \$71.02.¹¹ The average number of monthly visits in the first year (Column 2) is 4.68. The resulting price per average attendance (Column 3), \$15.15, is somewhat lower than for the monthly contract but still substantially higher than \$10. The lower estimate is presumably due to the selection of users with high expected attendance into the annual contract, and the lower fee of the annual contract. Table 4 shows the distribution across users of attendance and of the price per attendance in the first year of an annual membership. Only 24 percent pays less than \$10 per visit.

Stylized fact 1. (Price per expected attendance at enrollment) *Users who choose an unsubsidized flat-rate contract pay a price per average attendance of over \$17 in the monthly contract and over \$15 in the annual contract. The share of users who pay ex post less than \$10 per visit is 20 percent in the monthly contract and 24 percent in the annual contract.*

Robustness. Before we discuss possible explanations of Stylized Fact 1, we check its robustness and address concerns about institutional frictions.

1. *Sample.* So far we have restricted attention to the unsubsidized sample and pooled the results across clubs. We now consider subsidized users as well and disaggregate the results by club. The sample includes users initially enrolled with a monthly contract in the sample ‘First Spell’. Separately for each club, we regress health club attendance on the monthly fee using an Epanechnikov kernel. The measure of attendance is the average attendance per month in the first 6 months. We cross-validate club-by-club with a grid search to compute the optimal bandwidth for the price.¹² Figure 1a shows the results for club 1. The average monthly attendance from the kernel regression lies between 3 and 5 and is increasing in price, although the estimates are not very smooth given the small bandwidth suggested by the cross-validation. We use the average attendance from the kernel regression to compute the ratio of price and average attendance for each level of price. Figure 1b plots the price per average attendance with 95 percent confidence intervals. The price per average attendance is significantly higher than \$10 for users paying a monthly fee in excess of \$53. The estimates for club 2 are comparable (Figures 1c and 1d) and somewhat smoother given the larger optimal bandwidth. In club 3 the price per average attendance is higher than the per-visit fee of \$8 for users paying a fee in excess of \$46 (Figure 1f).

2. *Underrecording of attendance.* The high price per attendance could result from underrecording of attendance due to a faulty computer system or moral hazard problems with the staff. Alternatively, health club employees may simply seek to avoid queues of users waiting to swipe. While these phenomena may be observed in other health clubs, they are unlikely to occur at the three health clubs in our sample. These clubs put in place one of the most advanced

¹¹We use a daily discount factor of .9998, implying an adjustment factor $T(1 - \delta) / (1 - \delta^T)$ equal to 1.037.

¹²Pagan and Ullah (1999), pp. 110–120.

and reliable systems to track attendance in the industry, presumably in order to report attendances to the subsidizing corporations.¹³ Unlike in most clubs, a front-desk employee collects the cards from the members and swipes them in a quiet moment. Therefore, card swiping does not generate queues. Nevertheless, to assess the importance of occasional computer lapses or laziness of specific employees, we construct a test of accuracy of the attendance records. For each day, we calculate the fraction of members attending and regress it on a set of controls: 6 day-of-the-week dummies, 11 month dummies, 3 year dummies, and 15 holiday dummies. If recording of attendance is largely driven by random variation in computer crashes or staff dedication, the R^2 of this regression should be low. The R^2 of the regression for club 1 instead is as high as .8785, with the day-of-the-week dummies explaining most of the variance. The regression for clubs 2 and 3 yield an even higher R^2 of .8915.¹⁴ The high explanatory power of these regressions suggests that daily variation in recording precision is limited.

3. *Additional benefits.* The monthly and annual contracts provide the same benefits as the pay-per-visit system except for the option to rent an overnight locker at an extra fee.¹⁵ If users value this option highly, they may be willing to forgo the monetary savings of paying per visit. However, only 5.52 percent of the users ever rent a locker. If we exclude these users, the results on price per average attendance for the monthly contract do not vary.

4. *Ex-post subsidies.* Some HMOs reimburse members partially for health club expenses. To the extent that these reimbursements make the annual and the monthly contract cheaper relative to the pay-per-visit contract, they induce users to choose flat-rate contracts. However, a survey of the HMOs¹⁶ in the state where the three clubs operate shows that the discounts apply either only to the initiation fee, or both to flat-rate and pay-per-usage contracts.

Overall, we observe a robust deviation from the prediction of a model of time-consistent agents with rational expectations. Non-subsidized users enrolled in contracts with flat fees pay a price per average attendance that is significantly higher than the per-visit price available as an alternative contract. The result is robust to the type of contract (monthly or annual), the sample (the amount of subsidy), and the club considered. The results do not appear to depend on measurement error, unobserved benefits, or ex-post subsidies. The deviations from the predictions for time-consistent agents are large in size: unsubsidized members of a monthly contract pay 70 percent in excess of the \$10 fee.

¹³In fact, we selected these clubs in part because of the data quality. A dozen of other clubs with which we established preliminary contacts had software or hardware problems in the recording of attendance.

¹⁴Detailed results are available in DellaVigna and Malmendier (2003).

¹⁵In particular, a 10-visit card gives the same rights to get a towel and a temporary locker, hire a personal trainer, take the (free) aerobic classes and attend other clubs of the same company.

¹⁶We report the results in Appendix Table 3 in DellaVigna and Malmendier (2003). We thank Nancy Beaulieu for providing the list of HMOs.

3.3 Interpretation

The model in Section 3.1 suggests two possible explanations for this deviation based on time inconsistency and naiveté. First, sophisticated time-inconsistent users ($\beta = \hat{\beta} < 1$) may purchase the monthly contract as a commitment device to increase attendance. Under this interpretation, the inequality (4) and the empirical results imply a lower bound for $(1 - \beta)\delta b \cdot [G(\beta\delta b) - G(\beta\delta b - p)] / G(\beta\delta b)$ of $\$17 - \$10 = \$7$. In turn, this implies $(1 - \beta)\delta b \geq \7 . Second, users may be naive about their time inconsistency ($\beta < \hat{\beta} = 1$). Inequality (4) then implies that they overestimate the probability of attending the health club by at least seventy percent. An average attendance of two visits a week, while far from the actual number of visits, is a plausible estimate of the *desired* number of visits: the health club staff encourages members to attend two to three times per week.

We consider four alternative explanations within the framework of Subsection 3.1.

1. *Transaction costs.* Users may choose a flat-rate contract even though they expect to attend little if paying per visit entails large transaction costs. Formally, a per-visit transaction cost k adds to the per-visit price p in equation (5). Stylized fact 1 is consistent with the standard model for transaction costs of at least \$7 per visit. The actual transaction costs, however, appear to be small. Users can purchase a ten-visit pass by filling out a simple form, and can then enter the club for ten visits with the same procedure as users with a monthly or annual contract. A transaction-cost-based explanation requires a time cost of over \$70 for the few minutes necessary to fill out the form. A high distaste for payment per visit (Loewenstein and Prelec, 1998), however, could explain the findings.

2. *Risk aversion.* Assume a utility function that is additively separable in income and health club net benefits, $u(L, p) + v(c, b)$. Users that are risk averse in income may prefer a flat-rate contract to the pay-per-visit contract because the former contract minimizes the variance of the payments.¹⁷ This effect, however, should be small for the monthly contract. Over the small amounts of money required for a monthly contract, the agent is locally risk neutral (Rabin, 2001). The price per average attendance, instead, is particularly high for users in the monthly contract.

3. *Overestimation of net benefits.* Users may choose flat-rate contracts because they overestimate the future benefits of attendance b or underestimate the expected future costs $E[c]$, possibly because of projection bias (Loewenstein, O'Donoghue, and Rabin, 2002). This interpretation, like the one based on naiveté about self-control, suggests that agents overestimate future attendance due to non-rational expectations about an underlying parameter.

4. *Salesman techniques.* Given that users attend on average less than eight times per

¹⁷The result is not robust to the specification of the utility function. Under the assumption that the utility function is a concave function of the sum of income and health club net benefits, the predictions are reversed: more risk-averse agents are more likely to choose the pay-per-visit contract.

month, flat-rate contracts are more profitable for the health clubs than pay-per-visit contracts. Salespeople may pressure consumers to choose the monthly and annual contract. In the language of the model, this translates in an increase of the effective price p to reflect the psychic cost of disobedience to the salesman. In order to investigate whether consumers receive enough information about the pay-per-visit contract, we set up a small experiment. We provided subjects with incentives similar to those of a median user of health clubs. Each of the subjects had a budget of \$90 and had to choose the cheaper option to attend a club under the assumption that he/she would go on average four (three for some)¹⁸ times per month. The subjects could keep whatever they saved out of the \$90, in addition to a fixed payment of \$15. In order to ensure no communication with other subjects, we met each subject individually and at different times. We instructed them to visit club 1 in person, and met them again individually afterwards.¹⁹ Of the 11 subjects participating, 7 chose a pay-per-visit option (which was the pay-maximizing choice), while 4 picked the monthly contract. While the majority of subjects did not find it difficult to find out about the pay-per-visit option, it appears that some salespeople were reluctant to mention the pay-per-visit contract. As an alternative test, we consider the contractual choices of members of a specific HMO who can choose between a 20% discount on the flat-rate contracts and a \$6 payment per visit. Presumably, members claiming this discount are aware of both options, since both are explicitly listed on the HMO website. Nevertheless, the price per expected attendance over months 1 to 6 for HMO members enrolling with a monthly contract equals \$10.65 (s.e. 0.25), significantly higher than the \$6 price per visit.

While transaction costs and risk aversion are unlikely to explain our empirical findings, psychological transaction costs, overestimation of net benefits, and salesman techniques may contribute to the explanation of Stylized Fact 1.

3.4 Survey results

We complement the results from the data set of health club consumers with survey data from a sample of 48 randomly chosen health club users in California (details in Section 2). In the survey, we elicit the expectations of respondents about their own attendance in the next month, September.²⁰ This question attempts to measure directly whether health club users have rational expectations, as posited by the standard model and by the alternative models of sophisticated time inconsistency, transactions costs, and salesmen techniques. Although we do not observe actual attendance among these 48 survey respondents, it is unlikely to

¹⁸Interestingly, club 1 increased the price of a visit to \$20 and the price of a 10-visit-pass to \$150 after this paper was written. To make the choices of subjects comparable to the choices in our sample, we lowered the attendance to three visits per month.

¹⁹Instructions are available in DellaVigna and Malmendier (2003).

²⁰In our sample attendance in September is five percent lower than the yearly average.

differ substantially from attendance in our data set, which is very robust across demographic subgroups. Across 24 (gender)*(club)*(age) subgroups, the average monthly attendance over the membership is lower than 4.75 visits for 23 out of 24 groups, with an overall average of 4.17 monthly visits. We establish the following stylized fact.

Stylized Fact 2 (Forecasts of attendance). *The average forecasted number of monthly visits, 9.50 (s.e. 0.66), is more than twice as large as average attendance, 4.17.*

The overestimation displayed by the subjects matches with Stylized Fact 1. If health club consumers expected to attend 9 times per month, they should indeed choose a flat-rate contract, rather than paying per visit.

We also present the subjects with the following scenario: ‘Suppose that, based on your previous experience you expect to attend on average 5 times per month (about once a week), if you enroll in a monthly membership. You plan to attend the health club throughout the next year. Would you choose a monthly contract with a monthly fee of \$70 per month or 10-visit passes for \$100 (each visit costs \$10)?’ This question attempts to measure whether users endowed with rational expectations about attendance would still overwhelmingly choose flat-rate contracts. In the hypothetical scenario, 18 consumers out of 48 prefer the monthly contract, and 30 prefer the 10-visit pass. With realistic expectations about attendance, therefore, the majority of people prefers to pay per visit.

These two findings support the explanations of the results based on irrational expectations about self-control ($\hat{\beta} > \beta$) or about the net benefits of attendance (\hat{G} left-shift of G). However, one should take responses to hypothetical questions with caution, particularly because the survey sample differs from the health club sample.

4 Contract choice over time

4.1 Model

In the previous Section, we have used the differences between flat-rate and pay-per-visit contracts to draw inferences about consumer preferences and expectations. In this Section, we take advantage of two differences in the renewal procedure between the flat-rate contracts. First, the renewal default differs. The monthly contract is automatically renewed and requires a small effort—sending a letter, cancelling in person—in order to discontinue the membership. The annual contract automatically expires after 12 months, and cancellation requires no effort. Second, members of the monthly contract can cancel at any month, while members of the annual contract are committed for a year. We consider the direct implications of these contractual differences as well as the indirect implications for sorting between the two contracts. We evaluate the impact of these differences on cancellation lag, survival probabilities, and average attendance over time under the standard model and for the alternative models

presented in Sections 3.1 and 3.3.

Calibration. Consider two agents with identical preferences and cost realizations c . One has enrolled in the monthly contract, the other in the annual contract. At the end of the contractual period, each consumer can either renew with a monthly or annual contract, or switch to payment per visit (which is equivalent to dropping out). Denote with s the daily savings from switching to the pay-per-visit contract. Assume cancellation costs $k_M = \$15$ for the monthly contract and $k_A = \$0$ for the annual contract, and a daily discount factor δ of .9998 (corresponding to a yearly discount factor of .93). A rational time-consistent agent switches to payment per visit if the transaction costs are smaller than the future discounted gains from switching, i.e., if $k < \delta s / (1 - \delta)$ or $s > k(1 - \delta) / \delta$. For the calibrated values, the right-hand side equals approximately .3 cents under the monthly contract and 0 under the annual contract. Therefore, under either of the two contracts, a rational time-consistent agent switches at the next renewal whenever the savings are positive.

A time-inconsistent agent enrolled under the annual contract switches immediately to the pay-per-visit contract if the savings are positive ($s > 0$), given that the cancellation costs k_A are zero. The positive cancellation costs k_M of the monthly contract, instead, may induce a delay in cancellation since each self would like to delegate cancellation to a later self. The size of the delay depends on the sophistication of the agent.

A sophisticated time-inconsistent agent would like to delegate quitting to a later self, but prefers immediate quitting if cancellation would otherwise be postponed for too long. Following O'Donoghue and Rabin (2001), we obtain a bound on the delay in switching t_C , $t_C \lesssim k_M(1 - \beta\delta^{t_C}) / \beta\delta s$ or, under a calibrated²¹ β of .7, if $t_C \lesssim 6.6/s$. Thus, a sophisticated agent with savings of 20 cents per day ($s = .2$) or \$6 per month (one tenth of the average monthly fee) is willing to delay cancellation by at most 33 days. An agent who loses more is even less willing to delay. The cancellation costs for the monthly contract induce only a limited delay of cancellation. The probability of cancellation, therefore, is similar for the monthly and annual contract.

A naive time-inconsistent agent, instead, may delay switching forever in presence of the small cancellation costs k_M of the monthly contract. She believes that her future selves will be exponential. She thus incorrectly expects to switch in T periods if $k_M < \delta s / (1 - \delta)$. If this inequality holds, the agent postpones switching for T periods if $k_M \geq \beta\delta^T k_M + \beta\delta s(1 - \delta^T) / (1 - \delta)$, or, using the approximation $(1 - \delta^T) / (1 - \delta) \approx T$, if $k_M \gtrsim \beta\delta T s / (1 - \beta\delta^T)$. Once the T periods are over, however, the agent goes through the same reasoning, and postpones the decision for T more periods, and so on every T periods. If the agent believes that she can cancel

²¹Angeletos et al. (2001) and Paserman (2001) calibrate the hyperbolic model on field data and find values of β between .5 and .9. The results of the calibrations are similar for the ranges $k \in [\$5, \$20]$, $\delta^{365} \in [.90, .98]$, $\beta \in [.5, .9]$.

on any day²², that is, for $T = 1$, she delays cancellation forever is $s \lesssim k(1 - \beta\delta^T)/\beta\delta T \approx \6.6 or, equivalently, for monthly gains up to \$192. Differently from time-consistent and time-inconsistent sophisticated agents, naive agents thus may delay cancellation forever under the monthly contract, but not under the annual contract (which has no switching cost).²³

Of the alternative interpretations discussed in Sections 3.1 and 3.3, the ones based on transaction cost and overestimation of net benefits make similar predictions. Both forces make the flat-rate contracts more attractive relative to pay-per-visit contracts, and therefore lower s . These models therefore predict low cancellation rates under both the monthly and the annual contract, with no difference in cancellation rates between the two flat-rate contracts. Salesman pressure increases the renewal probability particularly for the monthly contract, in which one of the cancellation procedures involves contacting a salesman; under the annual contract, salesman pressure is likely to be less effective since cancellation is automatic.

Cancellation lags. The calibration results imply several predictions about cancellation behavior and contractual choice over time. A first, testable implication of the rational benchmark model applies to monthly members with low attendance. These consumers are likely to have high savings s from switching to payment per visit.

Prediction 3. (Cancellation lags) *Non-attenders cancel the monthly contract immediately.*

Sophisticated time-inconsistent consumers will accumulate only minimal delays. Prediction 3 does not apply for the other models, naiveté about self-control, high transaction costs of per-visit payment, overestimation of net benefits, or salesman techniques. In all of these cases, we may observe consumers with low attendance that nevertheless do not cancel the monthly contract.

Sorting. Consumers take the longer commitment of the annual contract into account when choosing their initial contract. Consumers who anticipate a high chance of being low-attenders in the future prefer the monthly contract. These users value highly the option to switch sooner to payment per visit. Users who, instead, believe that they will be high-attenders prefer the annual contract. These users value the reduced price of this membership and do not mind the yearly commitment.

The users selected into the annual contract, therefore, are more likely to be frequent users. Using attendance in the initial months (before the selective exit) as a measure of the likelihood to be a frequent user, we obtain the following prediction, which applies to any of the models in consideration.

²²The case $T = 1$ corresponds to a uniform prior over the monthly deadline, as suggested by the survey evidence presented in Section 2. If agents know the monthly deadline ($T = 30$), delay still occurs for substantial values of the savings variable s , that is, if the monthly gain from switching is at most \$6.6.

²³O'Donoghue and Rabin (2001) show that partially naive agents are similar to fully naive agents in their delaying behavior. We omit this case for brevity.

Prediction 4. (Sorting) *Average attendance in the initial months should be higher under the annual than under the monthly contract.*

Survival probability. The renewal behavior of monthly and annual users depends on the net effect of cancellation costs and sorting. We consider the contractual choice at a time when both contracts are up for renewal, i.e., after 12 or 24 months. We define the survival probability $S_{j,t}$ as the share of consumers who initially enrolled in contract j (equal to *Monthly* or *Annual*) and are still enrolled with any flat-rate contract after t months, with $t = 12, 24$. For example, $S_{M,12}$ is the share of monthly members who has not switched to payment per visit by month 12. Similarly, $S_{A,12}$ is the share of annual members who, after one year, renews with an annual or a monthly contract.

Sorting implies that the users selected into the annual contract are ex-post more likely to be frequent users, and therefore more likely to renew with a flat-rate contract. This increases $S_{A,t}$ relative to $S_{M,t}$. Cancellation costs for the monthly contract, instead, may act to increase $S_{M,t}$ relative to $S_{A,t}$. In the standard model with time-consistent agents with rational expectations, the calibrations above suggest that the direct effect of cancellation costs is very small. We therefore expect the sorting effect to dominate.

Prediction 5. (Survival probability) *The survival probability after one and after two years is higher for agents who initially chose the annual membership than for agents who initially chose the monthly membership: $S_{A,t} > S_{M,t}$, for $t = 12, 24$.*

We expect this prediction to hold for the cases of time inconsistency with sophistication, transaction costs, risk aversion, or overestimation of net benefits. Each of these modifications makes renewal more likely under both contracts, but not differentially so under the monthly contract.

We obtain a different prediction for two classes of models, time inconsistency with naiveté and models with salesman pressure. Small cancellation costs induce naive time inconsistent agents to delay cancellation under the monthly, but not under the annual contract. If the delay is strong enough to override sorting effects, the survival probability should be higher for the monthly contract: $S_{A,t} < S_{M,t}$ for $t = 12, 24$. Similarly, salesman techniques may lead to higher renewal under the monthly contract, since users can drop out of the annual contract without facing an health club employee.

While so far we have considered the unconditional survival probabilities S_A and S_M , the survival probabilities are functions of the realized cost type c . Although the cost type is unobserved, we can proxy for it with the attendance in the months prior to the 12th month.²⁴ We denote by $S_{j,t}(v)$ the survival probability of contract j at time t conditional on attendance v .

²⁴Formally, we assume that the distribution of past visits conditional on the costs c does not depend on the contract chosen. This condition is likely to be satisfied, since the monthly and the annual contract have the same (zero) per-visit price.

We expect that survival $S_{j,t}(v)$ should be increasing in past attendance under either contract. For agents with standard preferences and beliefs, we also expect higher survival rates under the annual than under the monthly contract for all levels of past attendance v (Figure 2a). Users with ex-ante low expected cost sort into the annual contract. Therefore, given that v is only a noisy proxy for the cost c , annual users are more likely to be low-cost types, and therefore to renew, than monthly user with the same attendance v .

Prediction 6. (Survival probability as a function of attendance) *The survival probability after one and after two years is higher for agents who initially chose the annual membership than for agents who initially chose the monthly membership at all levels of past visits v : $S_{A,t}(v) > S_{M,t}(v)$ for all v , $t = 12, 24$.*

Prediction 6 should hold also for the case of time inconsistency with sophistication, transaction costs, risk aversion, or overestimation of net benefits. For time-inconsistent agents with naive expectations, instead, the reverse inequality may hold— $S_{A,t}(v) < S_{M,t}(v)$ for all v , $t = 12, 24$ —given that the agents may delay cancellation forever under the monthly contract (Figure 2b.) Salesmen techniques may also induce a higher survival under the monthly contract for all levels of attendance v .

Finally, we can consider a setting where a fraction of the agents is naive and a fraction is not (Figure 2c). The survival probability for the heterogeneous population is a convex combination of the survival probabilities for agents with rational expectations (Figure 2a) and for naive agents (Figure 2b). If the proportion of naives is sufficiently large, we expect $S_{M,t}(v) \geq S_{A,t}(v)$ at least for low levels of attendance v . At these levels, the delay effect is stronger since more users have high costs of attendance and intend to quit. At high levels of attendance, few agents want to switch to the pay-per-visit contract and the sorting effect is likely to dominate.

Attendance over time. As time goes by, users learn about the effort of commuting to the club and the enjoyment of exercising. Only users who ex post attend frequently enough renew with a flat-rate contract. Learning therefore induces selective exit of individuals with ex-post low attendance patterns. Define as stayers individuals initially enrolled in an annual contract who do not switch to a pay-per-visit contract after the first year. Attendance for stayers in the second year should be higher than for the initial group of annual members in the first year, since the low-attenders have switched to paying per visit.²⁵ All the models considered above yield the following prediction.

Prediction 7. (Expected attendance over time for annual contract) *Among users initially enrolled in an annual contract, the expected attendance in the second year among stayers is higher than the expected attendance in the first year for the initial group.*

The standard model makes a parallel prediction for the monthly contract. As low-attenders

²⁵The pattern of average attendance within each year, instead, depends on the type of shocks.

quit, the average attendance for stayers increases from month to month, since users are allowed to quit in any month.

Prediction 8. (Expected attendance over time for monthly contract) *Among users initially enrolled in a monthly contract, the expected attendance among stayers should increase from month to month.*

The alternative models presented above all make this prediction, except for the model of naive agents and for salesman pressure. Naive agents delay cancellation in the monthly contract, even if they intend to quit. Since there is no selective exit, expected attendance among stayers need not increase over time. In fact, if negative shocks are more common than positive shocks²⁶, average attendance decreases over time. Similar conclusions follow if agents do not cancel because of salesmen pressuring them to renew.

4.2 Empirical Analysis

Cancellation lags. To test Prediction 3, we adopt, as a conservative measure of cancellation lag, the number of full months between the last attendance and contract termination for users who hold a monthly contract at the time of termination. For example, if an agent attends the last time on March 10 and cancels on April 5, we count the 51 days between last attendance (March 10) and membership termination (April 30) as *one* full month. We restrict the sample to users who paid no initiation fee.²⁷

Stylized fact 3. (Cancellation lags) *On average, 2.29 full months elapse between the last attendance and contract termination for monthly members, with associated membership payments of \$185. This lag is at least 4 months for 20 percent of the users.*

Even though the transaction costs of cancellation are likely lower than \$20 (time cost of sending a letter or visiting the club), users spend on average \$185 in membership fees after their last attendance.

Sorting. To test Prediction 4 on sorting, we compare the average number of visits in months 2, 3 and 4 of tenure for individuals initially enrolled in the monthly and in the annual contract.²⁸ Given that the price per visit p is zero for both contracts, differences in attendance should reflect differences in the expected attendance cost. Column 1 of Table 5 reports the results for the sample ‘First Spell.’ In each month, expected attendance is higher under the

²⁶This is the case if the process for the cost of attendance is mean reverting. The agents select into the flat-rate contract when they have a very low realization of costs which then reverts to the mean.

²⁷We include users with an unsubsidized membership (monthly fee higher than \$70 or annual fee higher than \$700) who joined the club before the month of April 1998.

²⁸We exclude the first month because attendance is pro-rated over the number of effective days of membership, and the pro-rating procedure is slightly different for the annual and the monthly contract. We do not extend the comparison to months after the fourth since users who experience a high cost can quit under the monthly contract but not under the annual contract.

annual than under the monthly contract, and significantly so in months 3 and 4. Overall, average attendance in months 2 to 4 is 9.5 percent higher under the annual contract. The magnitude of this difference is comparable to variation in average attendance by age groups and by gender. When we break down the sample into 24 age-gender-month cells, average attendance is higher under the annual contract in 20 cells out of 24. Even after controlling for some heterogeneity, individuals with lower cost of attendance are more likely to choose the annual contract at enrollment.

Stylized fact 4. (Sorting) *Average attendance in months 2-4 is 9.5 percent higher under the annual contract than under the monthly contract.*

Survival probability. We construct the survival measure s_i as follows. For spells starting with an annual contract, the survival of member i , s_i , equals 1 if no more than one calendar month elapses between the expiration of the first annual contract²⁹ and the enrollment of a new monthly or annual contract; s_i equals 0 otherwise. For spells starting with a monthly contract, no equally natural definition is available. We err on the side of overstating cancellation in the monthly contract and set s_i to 1 if the individual is enrolled on the 14th month of active, paid membership, and 0 otherwise.³⁰ We assume the following simple empirical specification to test Predictions 4 and 5:

$$s_i = 1 \text{ if } s_i^* = \alpha + \gamma M_i + \Phi V_i + \Pi(V_i * M_i) + BX_i + \varepsilon_i \geq 0, \quad (6)$$

where ε_i is normally distributed and M_i is a dummy variable that equals 1 if the first contract for individual i is a monthly contract, and 0 otherwise. The measure of attendance $V_i = [v_i, v_i^2, v_i^3, v_i^4]$ is a quartic in the average monthly attendance over all the available months until the 13th active month. Finally, the vector of controls X includes gender, a quadratic function of age, a dummy for corporate membership, a dummy for student membership, 11 dummies for the month and 4 dummies for the year of enrollment. We use the sample ‘First spell’ restricted to users who joined the club at least 14 active months before the end of the sample period. We also drop users with missing values of a control variable, as well as spells that are censored before the 14th active month.

Average survival probability. First, we constrain the coefficients Φ and Π to be zero. Within this specification, the coefficient γ captures the average difference in survival probability between users initially enrolled in a monthly contract and users initially enrolled in an annual contract. The coefficients in Table 6 are the marginal change in a coefficient in response to an infinitesimal change in the continuous independent variables, and a discrete change for the independent dummy variables. In the specification without controls (Column 1), $\hat{\gamma}$ is

²⁹In 11.5 percent of the cases, the first annual contracts lasts more than 12 months due to promotional months and freezing periods.

³⁰We exclude from the count of active months promotional periods, months in which the contract was frozen, and months (up to 3 in a row) in which the agent has temporarily quit the club.

positive. Enrollment in a monthly contract increases survival by 3.18 percentage point relative to the baseline rate of 39.93 percent survival with the annual contract. The introduction of demographic controls increases the coefficient $\hat{\gamma}$ from .0318 to .0509 (Column 2) and to .0514 (Column 3) with dummies for the time of enrollment. Controlling for some of the unobserved heterogeneity reduces the downward bias on the coefficient due to sorting and makes the coefficient significantly positive. For example, individuals enrolled with a monthly contract are significantly younger than users with an annual contract (Table 2), and young people are less likely to renew (Columns 2 and 3 of Table 6). Failing to control for age biases downward the coefficient $\hat{\gamma}$.

Stylized fact 5 (Survival probability). *The survival probability after 13 months for the monthly contract is 12.5 percent higher than for the annual contract.*

Survival probability as a function of attendance. In Columns 4-5 we allow for a linear effect of past attendance v ³¹. The parameter γ now captures the difference in survival probability between monthly and annual contract for low past attendance, i.e., for $v = 0$. The estimated $\hat{\gamma}$ is positive and very large, equal to .1650 without demographic controls (Column 4) and equal to .1803 with controls (Column 5). For this second specification, automatic renewal doubles the survival probability for low levels of attendance from the baseline level of 16.79 percentage points for the annual contract to 34.82 for the monthly contract. The coefficient on v indicates that one additional visit per month is associated with a 6.56 percentage points higher survival probability for users enrolled in an annual contract. The responsiveness of the survival probability to attendance for users under the monthly contract is about half the size.

Figure 2d plots the predicted survival probability for the more general specification with a quartic polynomial in past attendance (Column 6). For low levels of past attendance the monthly contract has a substantially higher survival probability. For higher levels of attendance, instead, the annual contract has a higher survival.

Stylized fact 6 (Survival probability as a function of attendance). *For low levels of past attendance, the survival probability under the monthly contract is twice as high as under the annual contract. For high past attendance, the annual contract has a higher survival.*

Robustness. In Table 7 we check the robustness of the findings. First, we replicate the results of Columns 4 and 5 of Table 6 using, as a proxy for attendance costs, the average monthly attendance in months 9 to 12, if available, or else in the last 4 months of membership³². The estimates are essentially unchanged. Second, we use the probability of enrollment in a monthly or annual contract at the 15th and at the 16th month after the joining date as alternative measures of survival.³³ With demographic controls, users initially enrolled in the

³¹We check that past attendance predicts future attendance on the first year of tenure for users in an annual contract, for whom selective exit is not an issue. A regression of attendance in the 12th month on attendance on each of the first 6 months gives an R^2 of .4024.

³²For users with spells shorter than four months, we use the attendance data for all the available months.

³³Measures of survival at earlier months are inappropriate. First, given the pro-rating of the first and last

monthly contract are 7.19 percentage points more likely to be enrolled at the 15th month (Column 4), and 5.82 percentage points more likely to be enrolled at the 16th month (Column 6) than users initially enrolled under the annual contract. Alternatively, we measure the survival after two years as the probability of enrollment at the 27th and 28th month after the joining date (Columns 7 through 10). In the specifications with controls, the estimate of γ is positive, although not significantly different from 0.

Third, we replicate the results of Columns 1 and 2 of Table 6 for the sample ‘First spell and no subsidy’ (Columns 11 and 12) and for the larger sample ‘First spell’ restricted to users who pay at least \$60 per month in the monthly contract or \$600 per year in the annual contract (Columns 13 and 14).³⁴ In the first, smaller sample the estimated $\hat{\gamma}$ has a similar magnitude as in the benchmark specification, but the estimates are imprecise. In the second, wider sample, the coefficient $\hat{\gamma}$ is positive and very large (.0925 with controls), as well as precisely estimated. Overall, the results on survival probability are robust to the measure of past attendance, the measure of survival, and the sample.

Attendance over time. Finally, we test Predictions 7 and 8 on the dynamics of average attendance. We first consider spells starting with an annual contract in the sample ‘First spell and no subsidy’ and lasting at least two years.³⁵ We display the results in Columns 1 to 3 of the bottom part of Table 8.

Stylized fact 7 (Average attendance over time in annual contract). *In the annual contract, average monthly attendance in the first year for the initial group, 4.69, is significantly lower than in the second year for stayers, 6.85.*

The difference in attendance between the two groups is large: the baseline group in the first year attends on average 46 percent less than the stayers. Consequently, the price per average attendance in the first year, \$15.15, is significantly higher than in the second year, \$10.77. The results for the larger sample ‘First spell’ are comparable (Columns 4 to 6 of Table 8).

Figure 3a shows the within-year dynamics of the price per average attendance. The sample at month t is given by the users in ‘First spell and no subsidy’ who have joined with an annual membership and are still enrolled with a flat-rate contract after t months of tenure. Over the first 12 months the price per average attendance increases from 12.5 to 18, as negative shocks accumulate. At renewal (months 12 and 13), the price per attendance is halved.

For spells starting with a monthly contract, the sample for average attendance at month t is given by the users in ‘First spell and no subsidy’ who have joined with a monthly membership month of an annual contract, an annual contract always extends until at least the 13th month. Second, about 10 percent of the annual spells lasts until the 14th month due to a free promotional month.

³⁴In both cases, we drop individuals who have missing values for a control or who joined the club later than 14 active months before the end of the sample period.

³⁵The results remain unchanged if we restrict the sample further to users who renew with an annual contract after 12 months.

and are still enrolled with a flat-rate contract after t months of tenure. Columns 1 to 3 of the top part of Table 8 show the results by six-month groups.

Stylized fact 8 (Average attendance over time in monthly contract). *Average monthly attendance in the first six months of a monthly contract, 4.85, is 26 percent higher than in the next six months and is significantly higher than in any of the later six-month periods among stayers.*

The price per average attendance for the first six months, \$17.13, is significantly lower than in any of the later six-month periods.³⁶ As Figure 3b shows, the price per average attendance increases over the first 10 months from about \$15 to about \$20, and remains constant thereafter. The results extend to the monthly spells in the sample ‘First spell’ (Columns 4 to 6).

4.3 Interpretations

Stylized Facts 4 and 7 confirm that standard economic forces such as sorting and learning are at work in the sample. Consumers sort into the monthly and annual contract based on the heterogeneous distribution of costs and benefits of attendance (Stylized Fact 4). Moreover, consumers that learn that they are low-attenders switch out of the annual contract, lowering the average attendance over time (Stylized Fact 7).

Stylized Facts 3, 5, 6, and 8, however, are hard to reconcile with the predictions of the standard theory. Unsubsidized monthly members spend on average \$185 for periods with no attendance before cancellation (Stylized Fact 3), despite small cancellation cost. In addition, after one year, more individuals are enrolled in the monthly contract, which allows more freedom to cancel, than in the annual contract (Stylized Fact 5). This result does not seem to arise because of sorting but despite sorting. The result is economically and statistically significant, robust across specifications, and highest for users with low attendance (Stylized Fact 6). Average attendance decreases by 26 percent between the first six months and the next six months in the monthly contract (Stylized Fact 8), a pattern contrary to the predictions of learning, and opposite to the one found for annual contracts (Stylized Fact 7).

A model of naive time-inconsistent consumers can explain these stylized facts. The combination of time inconsistency and naiveté generates a status-quo bias. Naive consumers are substantially more likely to renew if renewal is automatic than if renewal requires a minimal effort. Therefore, survival probability is higher under the monthly contract, which is automatically renewed, than under the annual contract (Stylized Facts 3, 5, and 6). Delay of cancellation by low-attendance users can also explain why average attendance increases over time in the monthly contract, but not in the annual contract (Stylized Facts 7 and 8).

Most of the alternative theories discussed in Sections 4.1 and 3.3 do not appear to explain all

³⁶The results remain unchanged if we restrict the sample further to users who have had a monthly contract at all times until month t .

of these facts. Models of sophisticated time-inconsistent agents, including other rational models of self-control such as the temptation model by Gul and Pesendorfer (2001), are unlikely to generate the patterns observed in the data, as the calibrations in Section 4.1 illustrate. Models of high transaction costs can explain Stylized Fact 3 on cancellation lags, but not Stylized Facts 5, 6, and 8. Agents that dislike paying per visit should be more likely to renew both monthly and annual contract. For these same reasons, the data does not support the hypothesis that agents renew the membership because they would regard cancellation as a personal failure (Gourville and Soman, 1998). This story also implies that users may switch from the monthly to the annual contract to signal a strong commitment to themselves. This switch instead happens for only 1.5 percent of the 7,079 spells initiated with a monthly contract.

A model with overestimation of future net benefits can explain the substantial cancellation lag of Stylized Fact 3. Once again, however, this hypothesis does not easily explain Stylized Facts 5, 6, and 8 on higher survival for the monthly than for the annual contract. An over-estimation story that can explain all the stylized facts requires that in addition consumers underestimate the future costs of cancellation (i.e., overestimate future efficiency). With this additional assumption, the overestimation model mirrors the model of time-inconsistency and naiveté. An advantage of the latter model is that it is a fully specified model that makes predictions about when the agents will underestimate and when they will overestimate. A second difference is that partially naive individuals look for commitment devices, while agents that overestimate their efficiency do not. Arguably, the demand for personal trainers indicates a demand for commitment.

An alternative model that can rationalize most of the facts is a version of the salesman model. Health club employees may exert pressure on monthly members that show up at the club intending to cancel, but are bound to be less effective with annual members that let their membership expire. This may explain Stylized Facts 3, 5, 6, and 8. However, if individuals are aware of their vulnerability to social pressure, the share of annual memberships should be higher than the observed 13 percent share (Table 2). In addition, this explanation does not address the survey evidence of overestimation of future attendance (Stylized Fact 9).

4.4 Heterogeneity

While the model of naive time-inconsistent agents fits well all qualitative features of the data, some features of the data are suggestive of heterogeneity in the population. In particular, the patterns in Figure 2d fit well a heterogeneous population with naive and standard agents (Figure 2c). This type of heterogeneity predicts a positive correlation between different correlates of naiveté. Members with a long cancellation lag in the monthly contract should be more likely to pay a high price per attendance in the period before the last attendance, since naiveté induces both behaviors.

We test this prediction for users enrolled in the monthly contract. As a measure of cancellation lags, we use the number of consecutive full months between the last attendance and the expiration (as in Section 4.2). As a measure of price per attendance, we take the ratio of the payments to the health club over the attendance for the period between sign-up and n months before the last attendance, with n equal to 1, 2, 3 and 4. We limit the time frame in order to avoid a spurious correlation between the price per attendance and months of delay due to low attendance in the final months. Finally, we take the log of 1 plus the measures in order to reduce the skewness of both variables. The correlation between the delay and the price per attendance is positive and significant, with values between .213 ($n = 1$) and .204 ($n = 4$). Longer lags n between the two measures do not affect the estimate, suggesting that the correlation is not likely to be spurious.

Stylized fact 9 (Correlations). *Users who pay a high price per attendance in the monthly contract display a longer gap between last attendance and contract termination.*

Similarly, we predict that individuals who accumulate a long delay in cancellation should be less likely to freeze a contract if they face a temporary period of non-attendance. We find a highly significant negative correlation of -.1035 between a freezing proxy³⁷ and the cancellation delay. These results are consistent with the idea that time inconsistency and naive expectations drive both the results on the high price per attendance for flat-rate memberships (Section 3.2) and the results on renewal behavior (Section 4.2).

5 Size of effects

The behavior of the average health club user deviates systematically from the predictions of the standard theory. We now provide a coarse measure of the monetary size of this deviation. For monthly and annual memberships, we compute the difference between actual expenses over the whole enrollment spell and imputed expenses for the same number of attendances with 10-visit passes.³⁸ A positive value of this ‘average loss’ measure indicates that the user would have saved money purchasing 10-visit passes, and a negative value indicates that the user would have lost money with these passes. We construct the average loss measure for the sample ‘First spell and no subsidy.’ To reduce the likelihood of censoring, we consider only spells that start before October 1997.

The average loss per spell (Column 1 of Table 9) is \$698 for agents initially enrolled in a

³⁷The raw measure of freezing is a dummy variable which equals 1 if the individual ever froze before the last attendance. To correct for periods of non-attendance, we run a probit of this dummy on the longest consecutive number of months with no attendance before the last attendance, and the number of periods longer than 2 months with no attendance. The residual of this regression is the final measure of freezing.

³⁸We neglect the fact that attendance would be lower under a pay-per-visit contract than under a flat-rate contract. For time-consistent agents, the measure we adopt understates the savings from paying per visit; for time-inconsistent agents, it may overstate the savings.

monthly contract. This amount corresponds to 47.87 percent of the \$1,517 spent on the health club membership during the whole spell. For agents initially enrolled in an annual contract, there is a small but insignificant gain of \$61 (Column 4). Columns 2 to 3 and 5 to 6 present the same variables for samples that include later spells. The estimates of the loss for the monthly contract are somewhat lower, given the higher percentage of spells censored. The estimates of loss for the annual contract are higher (a loss of about \$200) but still imprecise.

Overall, the observed deviations from the standard model have large monetary consequences for users in the monthly contract, for whom the automatic renewal amplifies the effects of the initial overestimation of attendance. For users in the annual contract, the automatic expiration moderates the possible losses from the initial overestimation.

6 Conclusion

Do consumers display time inconsistency and overconfidence in the market? In this paper we have used a new panel data set from three US health clubs to provide an answer to these questions. Members who choose a contract with a flat monthly fee of over \$70 attend on average 4.8 times per month. They pay a price per expected visit of more than \$17, even though a \$10-per-visit fee is also available. On average, these users forego savings of \$700 during their membership. We present additional results on the interval between last attendance and contract termination, the survival probability, the average attendance over time, and the correlation between different behaviors. These results are difficult to reconcile with the standard assumptions of time-consistent preferences and rational expectations. A model of agents with time inconsistent preferences and overconfidence about self-control explains the findings. The agents overestimate the future attendance and delay contract cancellation whenever renewal is automatic. The leading alternative explanations are based on social pressure exerted by the salesmen and overestimation of future efficiency.

The results in the paper have implications for the contract design by firms. Rational, profit-maximizing health clubs can easily learn the features of consumer behavior using data sets like the one analyzed in this paper. We therefore expect them to offer contracts that are designed to maximize profits given the features of consumers preferences and beliefs. In a related paper (DellaVigna and Malmendier, 2002), we characterize the features of the profit-maximizing contract for goods that challenge the self-control of consumers with time-inconsistent preferences. For goods with immediate costs and delayed benefits—such as health club attendance—the profit-maximizing contract involves below marginal cost pricing of attendance and automatic renewal with a cancellation cost. The typical contract of health clubs in the Boston area indeed has these features. The evidence on contractual design confirms the conclusions of the analysis of consumer behavior.

A Data Appendix

The data on consumer behavior come from the attendance panel and the billing records. A 7-digit identification number allows us to link multiple spells of the same individual.

Attendance panel. Each time a user with a flat-rate contract exercises, a staff member swipes the electronic card of the user, and therefore creates an attendance record. A line of the attendance panel consists of the individual id, the date of the visit, basic demographic information (birthday, gender), a code for short-term memberships, the enrollment and the expiration date (for members that terminated the membership). All the information but the date of the visits is constant for an individual.

Billing records. The health clubs keep an official record of the customer payments. The billing data provide detailed and accurate information about the category of users—retail (the default), student, family, corporate—as well as the type of transaction. Each line of the billing panel consists of the individual id, the date of the contractual transaction, the 4-digit code that identifies the transaction, and the price paid (if any). For example, line “1234567 1/1/98 R564 55” indicates that user 1234567 on Jan. 1, 1998 paid an out-of-pocket monthly fee of \$55 that applies to employees of company XYW. For the monthly contract, typical transactions are the payment of the initiation fee, of the monthly fee, or of additional items such as an overnight locker or a personal trainer. Other common codes involve monthly freezes of a membership, bounced payments, and termination of a membership for delinquency in the payments. For the annual contract, typical transactions are the payment of the initiation fee and of the annual fee.

We use the price stated in the records as a measure of the monetary payments to the clubs. We could alternatively use the 4-digit code and a conversion table (based on the prices as of August 2000) to recover an imputed price. The correlation between the two measures of price is .9668. None of the results changes if we use the imputed price instead of the actual price.

Monthly panel. We merge the attendance and the billing panel into a unique data set, and we then transform the data into a balanced panel with monthly observations. Each observation consists of a variable defining the membership (not enrolled/enrolled in a monthly contract/enrolled in an annual contract/in a freeze), the number of attendances in the month, and the price paid for the month. For an annual contract, the monthly price is 1/12th of the original price. In order to deal with monthly and annual contracts that start in the middle of a month, we pro-rate the fees for the first month. For the final month of an annual contract, we also pro-rate the fees. Monthly contracts always terminate on the last day of the month, so no pro-rating is needed for the last month.

Enrollment spells. We define an *enrollment spell* as a continuous temporal sequence of monthly or annual contracts, including possible freezes of the membership. If no more than one calendar month of non-enrollment separates two contracts, we still include them in one spell. For example, this is the case if an annual contract expiring on 15/1/98 is renewed on 17/3/98. The missing monthly payment may be due to an (unrecorded) one-month promotional offer, a delay in payment, or missing data for a monthly payment.

We consider an enrollment spell *censored* if either the spell is ongoing at the end of the panel or if the spell is followed by a short-term contract. One-month, two-month, three-month, and four-month contracts with automatic expiration are available, mostly for summer users. We do not analyze these relatively uncommon contracts, and therefore consider censored a spell that is followed by one of these contracts. We also consider a spell censored if it is followed by a sequence of months with no contract and attendance in at least half of the months. We assume that in these periods health club members are using a free temporary membership, which the clubs grant in various promotional or charitable initiatives.

References

- [1] **Akerlof, George A.** “Procrastination and Obedience.” *American Economic Review*, 81(2), May 1991, pp. 1–19.
- [2] **Angeletos, Marios; Laibson, David I.; Repetto, Andrea; Tobacman, Jeremy and Weinberg, Stephen.** “The Hyperbolic Buffer Stock Model: Calibration, Simulation, and Empirical Evaluation.” *Journal of Economic Perspectives*, Summer 2001, 15(3), pp. 47–68.
- [3] **Ariely, Dan, and Wertenbroch, Klaus.** “Procrastination, Deadlines, and Performance: Self-Control by Precommitment.” *Psychological Science*, May 2002, 13, pp. 219–224.
- [4] **DellaVigna, Stefano and Malmendier, Ulrike.** “Contract Design and Self-Control: Theory and Evidence,” 2002, mimeographed.
- [5] **DellaVigna, Stefano and Paserman M. Daniele.** “Job Search and Hyperbolic Discounting.” The Maurice Falk Institute for Economic Research in Israel (Jerusalem, Israel), Discussion Paper No. 00.15, December 2000.
- [6] **Fang, Hanming, and Silverman, Daniel.** “Time Inconsistency and Welfare Program Participation: Evidence from the NLSY”, mimeographed, 2001.
- [7] **Fielding, Helen.** *Bridget Jones’ Diary: A Novel*. Penguin USA, 1999.
- [8] **Fielding, Helen.** *Bridget Jones: The Edge of Reason*. Penguin USA, 2001.
- [9] **Gourville, John T. and Soman, Dilip.** “Payment Depreciation: The Behavioral Effect of temporally Separating Payments from Consumption” *Journal of Consumer Research*, September 1998, 25(2), pp. 160–174.
- [10] **Gruber, Jonathan and Koszegi, Botond.** “Is Addiction ‘Rational’? Theory and Evidence.” *Quarterly Journal of Economics*, November 2001, 116(4), pp. 1261–1303.
- [11] **Gruber, Jonathan and Mullainathan, Sendhil.** “Do Cigarette Taxes Make Smokers Happier?”, NBER working paper No. w8872, April 2002.
- [12] **Gul, Faruk and Pesendorfer, Wolfgang.** “Temptation and Self-Control.” *Econometrica*, November 2001, 69(6), pp. 1403–1436.
- [13] **Kirby, Kris N. and Herrnstein, R.J.** “Preference Reversals Due to Myopic Discounting of Delayed Reward.” *Psychological Science*, March 1995, 6(2), pp. 83–89.
- [14] **Laibson, David I.** “Golden Eggs and Hyperbolic Discounting.” *Quarterly Journal of Economics*, May 1997, 112 (2), pp. 443–77.
- [15] **Larwood, Laurie and Whittaker, William.** “Managerial myopia: self-serving biases in organizational planning.” *Journal of Applied Psychology* 62 (1977), pp. 194–198.
- [16] **Loewenstein, George and Prelec, Drazen.** “Anomalies in Intertemporal Choice: Evidence and an Interpretation.” *Quarterly Journal of Economics*, May 1992, 107(2), pp. 573–597.
- [17] **Loewenstein, George and Prelec, Drazen.** “The Red and the Black: Mental Accounting of Savings and Debt.” *Marketing Science*, 1998, 17, pp. 4–28.

- [18] **Loewenstein, George, O'Donoghue, Ted D. and Rabin, Matthew.** "Projection Bias in Predicting Future Utility", mimeographed, July 2002
- [19] **Madrian, Brigitte C. and Shea, Dennis.** "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior" *Quarterly Journal of Economics*, November 2001, 116(4), pp. 1149–1187.
- [20] **Malmendier, Ulrike and Tate, Geoffrey.** "CEO Overconfidence and Corporate Investment," 2002, mimeographed.
- [21] **Malmendier, Ulrike and Tate, Geoffrey.** "Who Does Mergers? CEO Overconfidence and the Market's Reaction," 2003, mimeographed.
- [22] **Miravete, Eugenio J.** "Choosing the Wrong Calling Plan? Ignorance and Learning" *American Economics Review*, forthcoming.
- [23] **O'Donoghue, Ted D. and Rabin, Matthew.** "Doing It Now or Later." *American Economics Review*, March 1999, 89(1), pp. 103-124.
- [24] _____. "Choice and Procrastination." *Quarterly Journal of Economics*, February 2001, 116(1), pp. 121–160.
- [25] **Pagan, Adrian and Ullah, Aman.** *Nonparametric Econometrics*. Cambridge University Press: Cambridge, 1999.
- [26] **Paserman, M. Daniele.** "Job Search and Hyperbolic Discounting: Structural Estimation and Policy Evaluation." mimeographed, Hebrew University of Jerusalem, August 2001.
- [27] **Phelps, Edmund S. and Pollak R.A.** "On Second-Best National Saving and Game-Equilibrium Growth." *Review of Economic Studies*, April 1968, 35(2), pp. 85–199.
- [28] **Rabin, Matthew.** "Risk Aversion and Expected-Utility Theory: A Calibration Theorem." *Econometrica*, September 2000, 68(5), pp. 1281-1292.
- [29] **Samuelson, William and Zeckhauser, Richard.** "Status Quo Bias in Decision Making." *Journal of Risk and Uncertainty*, 1988, 1(1), pp. 7–59.
- [30] **Strotz, Robert H.** "Myopia and Inconsistency in Dynamic Utility Maximization." *Review of Economic Studies*, 1956, 23(3), pp. 165–180.
- [31] **Svenson, Ola.** "Are we all less risky and more skillful than our fellow drivers?" *Acta Psychologica* 47 (1981), pp. 143-148.
- [32] **Thaler, Richard H.** "Some Empirical Evidence on Dynamic Inconsistency." *Economics Letters*, 1981, 8, pp. 201–207.

Figure 1. Average attendance and price per average attendance (Kernel regressions)

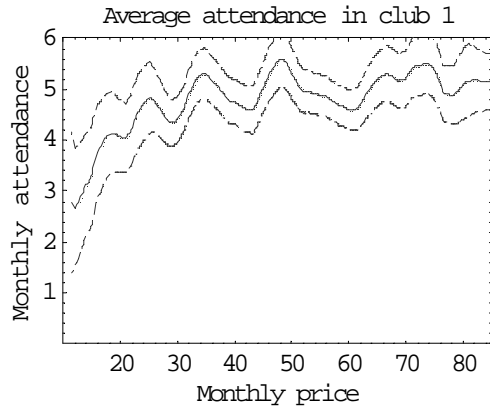


Figure 1a. Kernel regression of attendance on price (club 1, bandwidth 4).

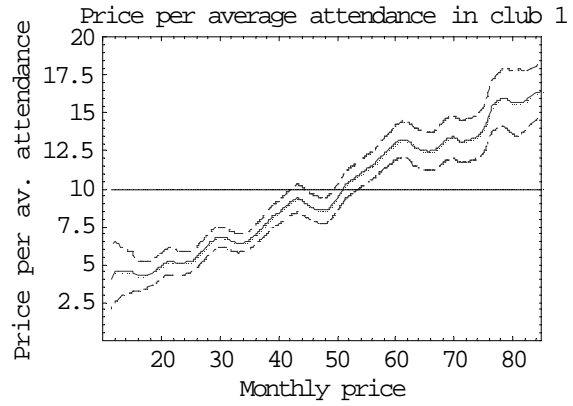


Figure 1b. Price per average attendance as a function of the monthly price (club 1, bandwidth 4).

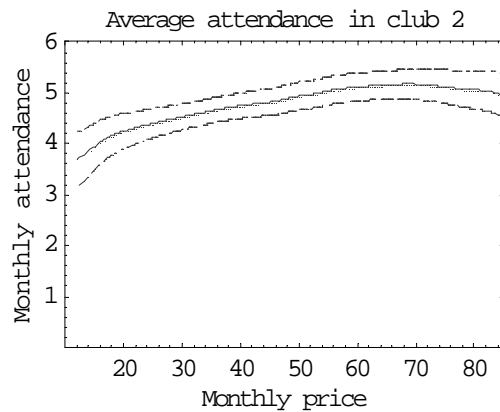


Figure 1c. Kernel regression of attendance on monthly price (club 2, bandwidth 16).

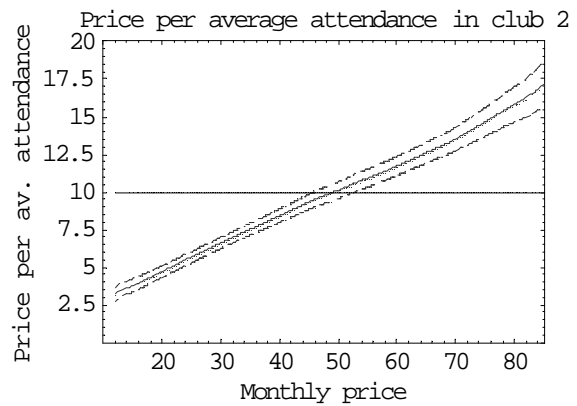


Figure 1d. Price per average attendance as a function of the monthly price (club 2, bandwidth 16).

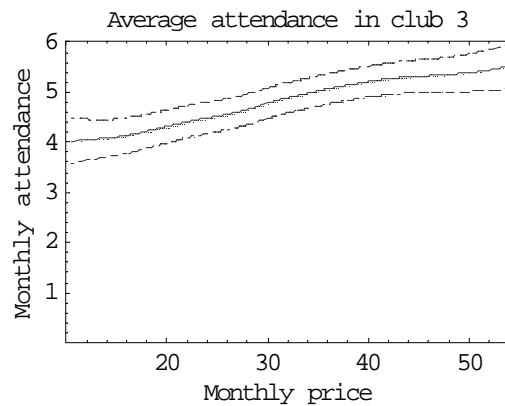


Figure 1e. Kernel regression of attendance on price (club 3, bandwidth 16)

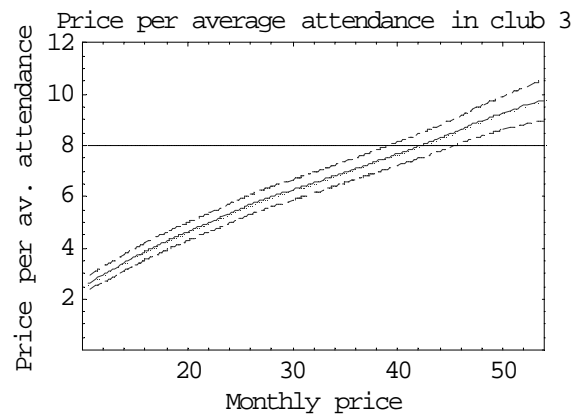


Figure 1f. Price per average attendance as a function of the monthly price (club 3, bandwidth 16).

Notes: Point estimates and 95 percent confidence intervals plotted. The sample is 'First spell' for individuals initially enrolled with a monthly contract. The individual price variable is the average price over the first six months. The individual attendance variable is the average attendance over the first six months. Figures 1a, 1c, and 1e show a kernel regression of attendance on price using an Epanechnikov kernel. The bandwidth is determined by cross-validation with a grid search separately for each club. Figures 1b, 1d, and 1f show the ratio of the price and the expected attendance predicted for that price using the kernel regression. Confidence intervals are derived using the Delta method.

Figure 2. Survival probability (as a function of attendance)

Figure 2.a.
Time-consistent or sophisticated agents

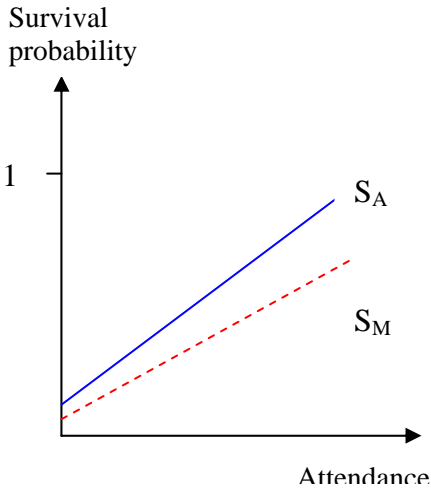


Figure 2.b.
Naïve agents

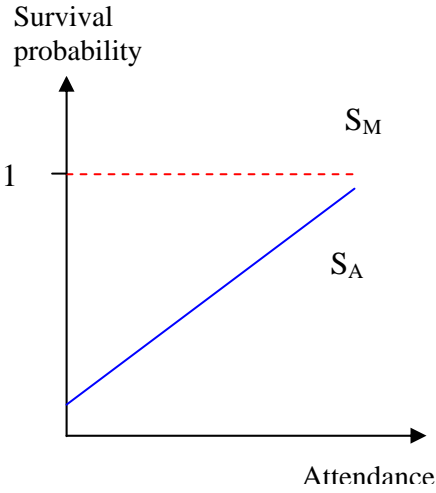


Figure 2.c.
Heterogeneity (time-cons./soph. and naïve agents)

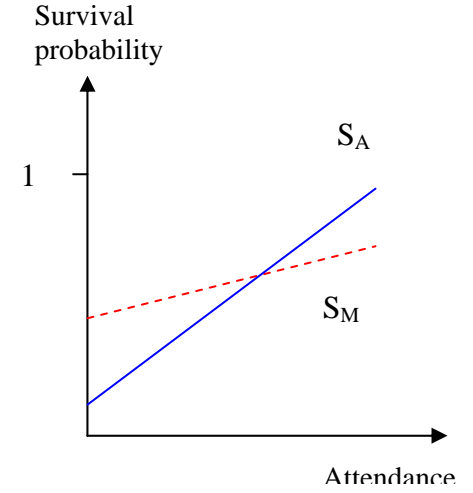
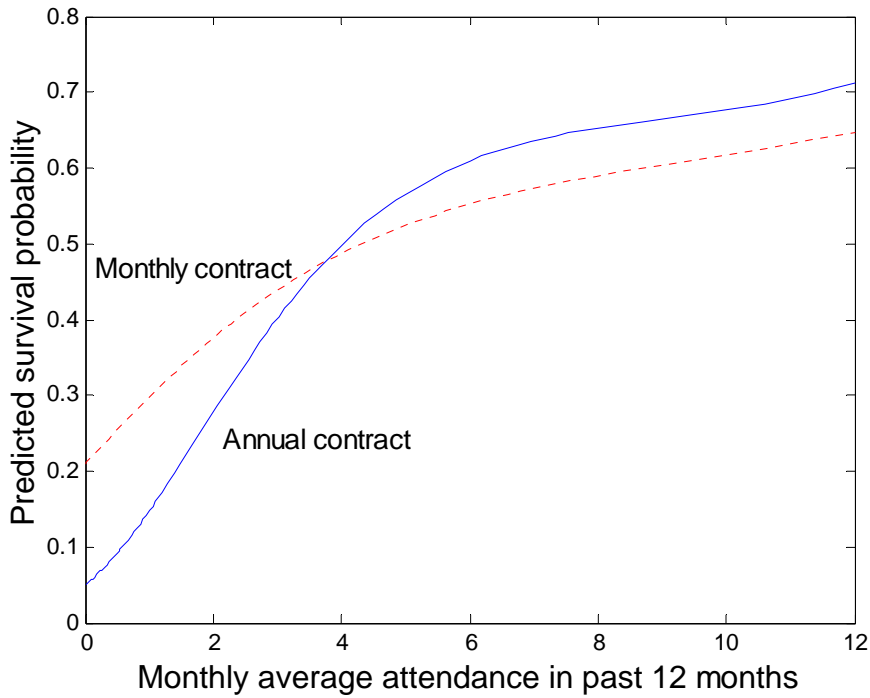


Figure 2.d. Predicted survival probability.
Probit specification with quartic polynomial in past attendance



Notes: Figure 2d plots the predicted probability of membership at the 14th active month after enrollment as a function of a quartic polynomial in past attendance. Figure 2d follows the probit specification of Column 6 in Table 6.

Figure 3. Price per average attendance over time

Figure 3a. Price per average attendance
Annual contracts with annual fee \geq \$700

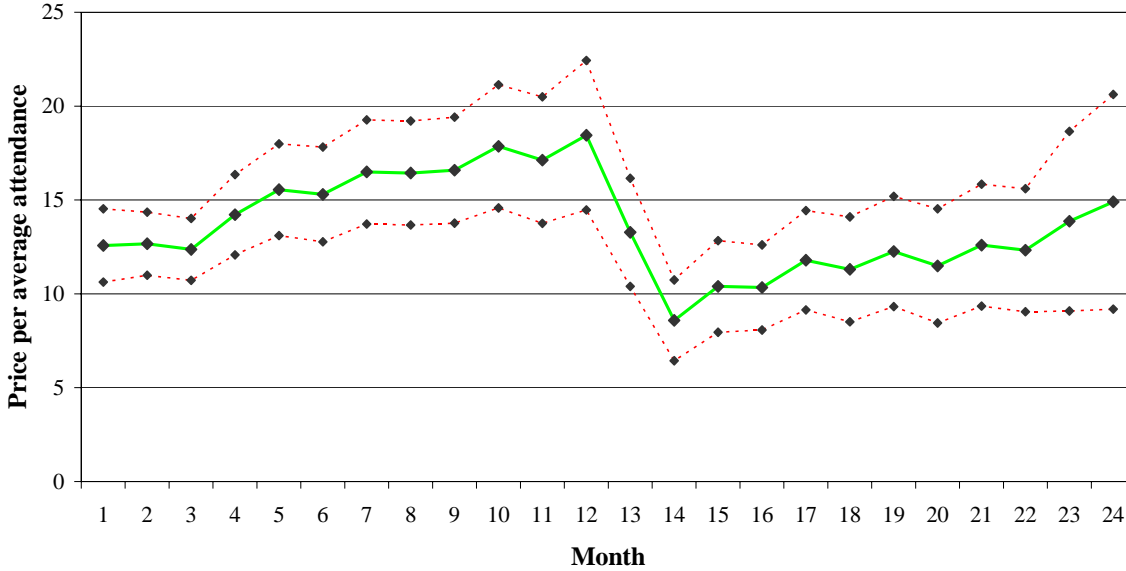
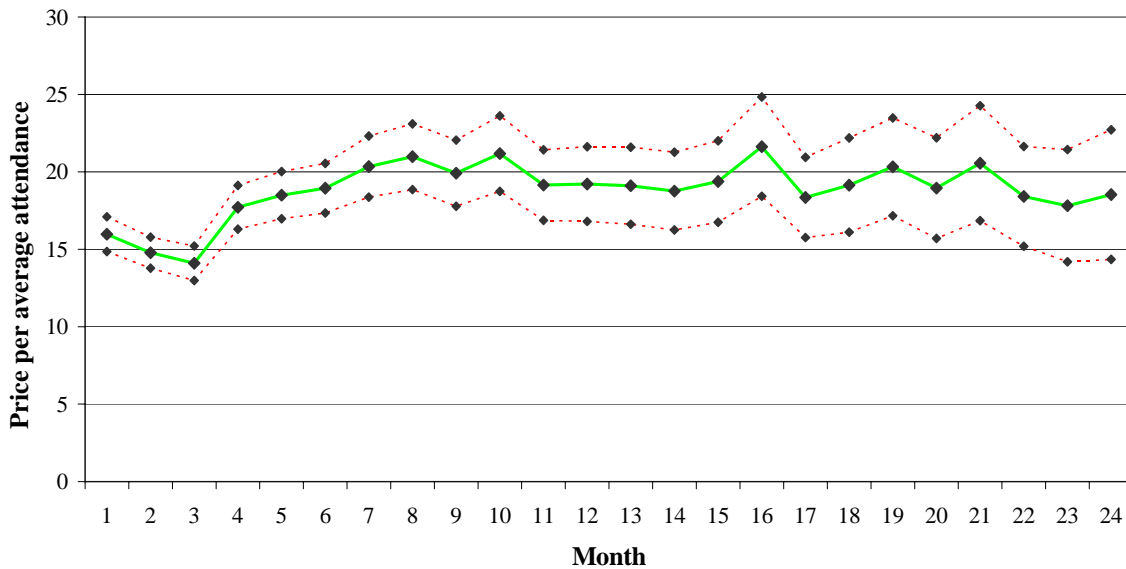


Figure 3b. Price per average attendance
Monthly contracts with monthly fee \geq \$70



Notes: Point estimates and 95 percent confidence intervals plotted. Figure 3a plots the ratio of average price and average attendance at month n of tenure. The sample is 'First spell and no subsidy, all clubs' for individuals initially enrolled in the annual contract and still enrolled at month n of tenure. Figure 3b plots the ratio of average price and average attendance at month n of tenure. The sample is 'First spell and no subsidy, all clubs' for individuals initially enrolled in the monthly contract and still enrolled at month n of tenure. Standard errors for the ratio of average price and average attendance computed using the bivariate Delta method.

Table 1: Stylized Facts and Explanations

Time-consistent agents (1)	Sophisticated time-inconsistent agents (2)	Partially naive time-inconsistent agents (3)	Trans. costs of payment per usage (4)	Overestimation of net benefits (5)	Salesman techniques (6)
Stylized fact 1. Price per average attendance > \$10	commitment	commitment, overestimation of attendance	distaste of paym. per usage	overestimation of attendance	pressure of salesman
Stylized fact 2. Forecast of 9.5 monthly visits compared to 4.2 actual monthly visits		overestimation of attendance		overestimation of attendance	
Stylized fact 3. Interval between last attendance and termination 2.3 full months		delay in cancellation	distaste of paym. per usage	overestimation of attendance	pressure of salesman
Stylized fact 4. Average attendance in first 4 months higher in annual than monthly contract	sorting	sorting	sorting	sorting	sorting
Stylized fact 5. Survival probability at 14th month 12.5 percent higher for monthly than for annual contract		delay in cancellation			pressure of salesman
Stylized fact 6. Survival probability at 14th month double for monthly than for annual contract for low past attendance		delay in cancellation			pressure of salesman
Stylized fact 7. Average attendance 46 percent higher in second year for annual contract	learning	learning	learning	learning	learning
Stylized fact 8. Decreasing average attendance over time in monthly contract		delay in cancellation			pressure of salesman
Stylized fact 9. Positive correlation of price per average attendance and interval between last attendance and termination		heterogeneity in naiveté			

Table 2: Descriptive Statistics

	Sample: First spell				Sample: First spell		Sample: First spell and no subsidy	
	Club 1	Club 2	Club 3	All clubs	All clubs		All clubs	
	All Contr. (1)	All Contr. (2)	All Contr. (3)	All Contr. (4)	First Contract Monthly (5)	First Contract Annual (6)	First Contract Monthly (7)	First Contract Annual (8)
Number of spells								
total	3548	2984	1446	7978	7079	899	912	208
completed spells	2440	1850	994	5284	4775	509	588	111
Total Amount in \$	569.05 (509.94) N = 3548	548.97 (559.85) N = 2984	312.24 (307.50) N = 1446	515.00 (509.07) N = 7978	501.28 (513.41) N = 7079	622.97 (459.66) N = 899	920.51 (713.68) N = 912	1041.80 (543.94) N = 208
Initiation fee	6.32 (26.64) N = 3548	1.99 (12.16) N = 2984	2.85 (12.88) N = 1446	4.07 (20.13) N = 7978	3.87 (19.45) N = 7079	5.62 (24.82) N = 899	14.46 (41.66) N = 912	17.07 (45.15) N = 208
Average fee per month								
monthly contract	52.40 (18.34) N = 3185	49.37 (18.93) N = 2663	31.36 (10.91) N = 1314	47.42 (19.10) N = 7162	47.31 (19.05) N = 7079	56.50 (20.51) N = 83	78.52 (5.07) N = 912	74.82 (15.33) N = 22
annual contract	48.33 (18.35) N = 445	43.76 (17.34) N = 405	24.15 (8.62) N = 151	42.83 (17.50) N = 1001	45.94 (15.77) N = 102	42.48 (17.66) N = 899	69.89 (4.19) N = 7	66.26 (4.21) N = 208
Average attendance per month								
monthly contract	4.06 (3.85) N = 3185	4.02 (3.82) N = 2663	3.75 (3.67) N = 1314	3.99 (3.81) N = 7162	3.98 (3.81) N = 7079	4.62 (3.79) N = 83	3.96 (3.77) N = 912	5.45 (4.15) N = 22
annual contract	4.45 (3.90) N = 445	4.22 (4.06) N = 405	4.16 (3.98) N = 151	4.31 (3.97) N = 1001	5.76 (4.20) N = 102	4.15 (3.92) N = 899	6.07 (4.04) N = 7	4.26 (3.87) N = 208
Contract choice per spell								
months with monthly contract	9.18 (8.42) N = 3548	8.91 (9.14) N = 2984	8.86 (8.91) N = 1446	9.02 (8.78) N = 7978	10.11 (8.70) N = 7079	0.44 (2.12) N = 899	11.70 (9.06) N = 912	0.55 (2.35) N = 208
months with annual contract	1.58 (4.75) N = 3548	1.95 (5.79) N = 2984	1.41 (4.84) N = 1446	1.69 (5.18) N = 7978	0.15 (1.52) N = 7079	13.76 (7.47) N = 899	0.08 (1.08) N = 912	15.16 (7.97) N = 208
freezing	0.28 (0.97) N = 3548	0.30 (1.12) N = 2984	0.18 (0.73) N = 1446	0.27 (1.00) N = 7978	0.30 (1.05) N = 7079	0.05 (0.37) N = 899	0.37 (1.22) N = 912	0.04 (0.32) N = 208
Female	0.44 (0.50) N = 3539	0.49 (0.50) N = 2984	0.47 (0.50) N = 1446	0.46 (0.50) N = 7969	0.48 (0.50) N = 7071	0.34 (0.47) N = 898	0.38 (0.49) N = 912	0.35 (0.48) N = 208
Age at sign-up	30.71 (8.43) N = 3343	31.54 (8.94) N = 2855	35.05 (9.27) N = 1363	31.81 (8.92) N = 7561	31.52 (8.78) N = 6710	34.04 (9.65) N = 851	33.14 (9.70) N = 855	34.40 (10.78) N = 197
Corporate member	0.43 (0.50) N = 3548	0.61 (0.49) N = 2984	0.43 (0.50) N = 1446	0.50 (0.50) N = 7978	0.50 (0.50) N = 7079	0.53 (0.50) N = 899	0.16 (0.37) N = 912	0.16 (0.37) N = 208
Student	0.05 (0.22) N = 3548	0.00 (0.04) N = 2984	0.00 (0.06) N = 1446	0.02 (0.15) N = 7978	0.02 (0.15) N = 7079	0.01 (0.12) N = 899	0.00 (0.06) N = 912	0.00 (0.07) N = 208

Notes: Standard deviation in parentheses. An enrollment spell starts whenever an individual enrolls (or re-enrolls) in the club and ends whenever the individual quits or is censored. The sample "First spell" consists of the first enrollment spell. The sample "First spell and no subsidy" restricts the sample "First spell" to those spells in which the average adjusted monthly fee is at least \$70 if the spell starts with a monthly contract and at least \$58 if the spell starts with an annual contract. The spells in column "First Contract Monthly" start with a monthly contract. The spells in column "First Contract Annual" start with an annual contract. "Average price per month" refers to the out-of-pocket fee in the case of corporate users.

Table 3: Price per Average Attendance at Enrollment

Sample: First spell and no subsidy, all clubs			
	Average price per month (1)	Average attendance per month (2)	Average price per average attendance (3)
Users initially enrolled with a monthly contract			
Month 1	55.09 (0.78) N = 873	3.45 (0.13) N = 873	15.98 (0.57) N = 873
Month 2	80.53 (0.44) N = 797	5.45 (0.18) N = 797	14.78 (0.51) N = 797
Month 3	70.02 (1.04) N = 780	4.97 (0.18) N = 780	14.09 (0.57) N = 780
Month 4	81.72 (0.26) N = 766	4.61 (0.19) N = 766	17.71 (0.72) N = 766
Month 5	81.87 (0.25) N = 701	4.43 (0.18) N = 701	18.50 (0.78) N = 701
Month 6	81.88 (0.28) N = 639	4.32 (0.19) N = 639	18.94 (0.82) N = 639
Months 1 to 6	83.00 (0.40) N = 912	4.85 (0.14) N = 912	17.13 (0.52) N = 912
Users initially enrolled with an annual contract, join 14 month before the end of sample period.			
Year 1	71.02 0.50 N = 145	4.69 0.38 N = 145	15.15 1.24 N = 145

Notes: Standard errors in parentheses. Standard errors for "Average price per average attendance" measure computed using the bivariate Delta method. The number of observations is denoted by N. An enrollment spell starts whenever an individual enrolls (or re-enrolls) in the club and ends whenever the individual quits or is censored. The sample "First spell" consists of the first enrollment spell. The sample "First spell and no subsidy" restricts the sample "First spell" to those spells in which the average adjusted monthly fee is at least \$70 if the spell starts with a monthly contract and at least \$58 if the spell starts with an annual contract. The sample for the t -th month includes spells that are ongoing, not frozen, and not miscoded at month t . For the 6-month period, the sample includes spells that are ongoing, not frozen, and not miscoded in at least one month in the period. For the 1-year period in the annual contract, the sample includes only spells that started at least 14 months before the end of the sample period, and that were not prematurely terminated because of medical reasons or relocation.

The "Average price" in period t is the average fee across people enrolled in period t . The "Average attendance" in period t is the average number of visits across people enrolled in period t . The measure in Column (3) is the ratio of the measure in Column (1) and the measure in Column (2).

Table 4: Distribution of Attendance and Price per Attendance at Enrollment

Sample: First spell and no subsidy, all clubs				
	First contract monthly, Months 1-6 (monthly fee >= \$70)		First contract annual, Year 1 (annual fee >= \$700)	
	Average attendance per month (1)	Price per attendance (2)	Average attendance per month (3)	Price per attendance (4)
Distribution of measures				
10th percentile	0.33	7.59	0.23	6.05
20th percentile	1.00	10.17	0.86	8.67
25th percentile	1.50	11.39	1.23	10.65
Median	3.91	20.89	3.58	20.34
75th percentile	7.00	58.39	6.58	59.82
90th percentile	10.75	107.50	11.00	119.64
95th percentile	12.83	170.00	13.25	239.28
	N = 912	N = 912	N = 145	N = 145

Notes: The number of observations is denoted by N. An enrollment spell starts whenever an individual enrolls (or re-enrolls) in the club and ends whenever the individual quits or is censored. The sample "First spell" consists of the first enrollment spell. The sample "First spell and no subsidy" restricts the sample "First spell" to those spells in which the average adjusted monthly fee is at least \$70 if the spell starts with a monthly contract and at least \$58 if the spell starts with an annual contract. The spells in column "First Contract Monthly, months 1-6" start with a monthly contract. The spells in column "First Contract Annual, year 1" start with an annual contract. The variable "Price per attendance" is defined as the ratio of the average price over the average attendance over the first period (6 months for the monthly contract, one year for the annual contract).

Table 5: Average Attendance in Monthly and Annual Contracts (Sorting)

Average attendance during the n-th month since enrollment			
Sample: First spell, all clubs			
	Month 2	Month 3	Month 4
Monthly contract	5.500 (0.0658) N = 6380	4.998 0.069 N = 5783	4.592 0.070 N = 5390
Annual contract	5.797 (0.1865) N = 874	5.583 0.191 N = 858	5.151 0.188 N = 839

Notes: Standard errors in parentheses. The number of observations is denoted by N. An enrollment spell starts whenever an individual enrolls (or re-enrolls) in the club and ends whenever the individual quits or is censored. The sample "First spell" consists of the first enrollment spell. The spells in row "Monthly Contract" start with a monthly contract. The spells in row "Annual Contract" start with an annual contract. The sample in month n includes spells that are ongoing, not frozen, and not miscoded.

Table 6: Probit of Renewal Decision I

Dependent variable: Enrollment at 14th active month							
Sample: First spell with non-missing controls, all clubs							
Controls:	no controls	controls	controls + time dummies	no controls	controls + time dummies	no controls	controls + time dummies
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dummy for enrollment with monthly contract	0.0318 (0.0217)	0.0509 (0.0217)	0.0514 (0.0218)	0.1650 (0.0321)	0.1803 (0.0317)	0.2858 (0.0508)	0.2943 (0.0502)
Average monthly attendance in the first 13 active months							
Attendance				0.0641 (0.0060)	0.0656 (0.0061)	0.2643 (0.0635)	0.2742 (0.0654)
(Attendance) ²						-0.0322 (0.0163)	-0.0339 (0.0170)
(Attendance) ³						0.0017 (0.0015)	0.0018 (0.0016)
(Attendance) ⁴						-0.00003 -0.00004	-0.00003 -0.00005
Monthly contract* (Average monthly attendance in the first 13 active months)							
Monthly*Attendance				-0.0292 (0.0063)	-0.0291 (0.0065)	-0.1429 (0.0654)	-0.1363 (0.0673)
Monthly*(Attendance) ²						0.0176 (0.0167)	0.0163 (0.0174)
Monthly*(Attendance) ³						-0.0008 (0.0015)	-0.0007 (0.0016)
Monthly*(Attendance) ⁴						0.00001 -0.00004	0.00001 -0.00005
Female		-0.0576 (0.0143)	-0.0566 (0.0144)		-0.0453 (0.0148)		-0.0458 (0.0149)
Age		0.0202 (0.0047)	0.0204 (0.0047)		0.0270 (0.0050)		0.0280 (0.0050)
Age square		-0.0002 (0.0001)	-0.0002 (0.0001)		-0.0003 (0.0001)		-0.0003 (0.0001)
Corporate member		0.0911 (0.0143)	0.0816 (0.0144)		0.1105 (0.0147)		0.1089 (0.0149)
Student member		-0.1342 (0.0500)	-0.1370 (0.0498)		-0.1071 (0.0530)		-0.0907 (0.0545)
Month and year of enrollment			X		X		X
Baseline renewal probability for monthly=0	0.3993	0.4033	0.4161				
Baseline renewal probability for monthly=0 and attendance=0				0.1598	0.1679	0.0497	0.0497
Number of observations	N=4905	N=4905	N=4905	N=4905	N=4905	N=4905	N=4905

Notes: Standard errors in parentheses. The number of observations is denoted by N. Entries in the Table represent the marginal coefficients of the probit in response to an infinitesimal change in the continuous variables, and a discrete change for the dummy variables. An enrollment spell starts whenever an individual enrolls (or re-enrolls) in the club and ends whenever the individual quits or is censored. The sample "First spell with non-missing controls" consists of the first enrollment spell for individuals for whom the demographic controls "age" and "female" are available. The sample is further restricted to individuals who join at least 14 active months before the end of the sample period. See the text for a definition of Enrollment at the 14th active month. The controls "Month and year of enrollment" indicate that the probit contains 11 dummies for the month of enrollment and 4 dummies for year of enrollment.

Table 7: Probit of Renewal Decision II. Robustness

Sample:	First spell with non-missing controls, all clubs										No subsidy I, all clubs		No subsidy II, all clubs	
	Enrollment at 14th active month		Enrollment at the 15th month		Enrollment at the 16th month		Enrollment at the 27th month		Enrollment at the 28th month		Enrollment at 14th active month		Enrollment at 14th active month	
	No Controls	Controls + Time Dummies	No Controls	Controls + Time Dummies	No Controls	Controls + Time Dummies	No Controls	Controls + Time Dummies	No Controls	Controls + Time Dummies	No Controls	Controls + Time Dummies	No Controls	Controls + Time Dummies
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
Dependent variable:														
Controls:														
Dummy for enrollment with monthly contract	0.1545 (0.0287)	0.1698 (0.0286)	0.0543 (0.0217)	0.0719 (0.0219)	0.0376 (0.0221)	0.0582 (0.0222)	-0.0009 (0.0261)	0.0262 (0.0252)	-0.0016 (0.0264)	0.0294 (0.0252)	0.0527 (0.0479)	0.0465 (0.0501)	0.0812 (0.0370)	0.0925 (0.0378)
Average attendance in last 4 months	0.0728 (0.0062)	0.0736 (0.0063)												
Monthly contract* (Average attendance in last 4 months)	-0.0384 (0.0065)	-0.0379 (0.0066)												
Female		-0.0462 (0.0149)		-0.0405 (0.0143)		-0.0405 (0.0144)		-0.0759 (0.0166)		-0.0811 (0.0167)		-0.0443 (0.0398)		-0.0306 (0.0280)
Age		0.0256 (0.0050)		0.0145 (0.0046)		0.0164 (0.0047)		0.0239 (0.0054)		0.0265 (0.0054)		0.0312 (0.0115)		0.0253 (0.0083)
Age square		-0.0003 (0.0001)		-0.0001 (0.0001)		-0.0002 (0.0001)		-0.0002 (0.0001)		-0.0003 (0.0001)		-0.0004 (0.0001)		-0.0003 (0.0001)
Corporate member		0.1004 (0.0149)		0.0747 (0.0144)		0.0700 (0.0145)		0.0713 (0.0168)		0.0705 (0.0169)		0.2071 (0.0484)		0.0043 (0.0319)
Student member		-0.1147 (0.0525)		-0.1151 (0.0503)		-0.0953 (0.0516)		-0.0916 (0.0591)		-0.0767 (0.0604)		0.1527 (0.2722)		-0.1616 (0.0669)
Month and year of enrollment		X		X		X		X		X		X		X
Renewal probability for monthly=0			0.3983	0.4162	0.3925	0.4077	0.2677	0.2853	0.2589	0.2730	0.4701	0.5426	0.4366	0.4373
Renewal probability for monthly=0 & attend.=0	0.1823	0.1832												
Number of observations	N=4905	N=4905	N=4990	N=4990	N=4860	N=4860	N=2874	N=2874	N=2777	N=2777	N=704	N=704	N=1362	N=1362

Notes: Standard errors in parentheses. The number of observations is denoted by N. Entries in the Table represent the marginal coefficients of the probit in response to an infinitesimal change in the continuous variables, and a discrete change for the dummy variables. An enrollment spell starts whenever an individual enrolls (or re-enrolls) in the club and ends whenever the individual quits or is censored. The sample "First spell with non-missing controls" consists of the first enrollment spell for individuals for whom the demographic controls "age" and "female" are available. The sample is further restricted to individuals who join at least 14 active months before the end of the sample period. The sample "No Subsidy I" is a restriction of the sample "First spell with non-missing controls" to individuals paying on average a per-month fee of at least \$70. The sample "No Subsidy II" is a restriction of the sample "First spell with non-missing controls" to individuals paying on average a per-month fee of at least \$60. See the text for a definition of Enrollment after 13 active months. The controls "Month and year of enrollment" indicate that the probit contains 11 dummies for the month of enrollment and 4 dummies for year of enrollment.

Table 8: Attendance and Price per Average Attendance Over Time

	Sample: First spell and no subsidy, all clubs			Sample: First spell, all clubs		
	Average price per month	Average attendance per month	Average price per average attendance	Average price per month	Average attendance per month	Average price per average attendance
	(1)	(2)	(3)	(4)	(5)	(6)
Users initially enrolled with a monthly contract						
Months 1-6	83.00 (0.40) N = 912	4.85 (0.14) N = 912	17.13 (0.52) N = 912	50.44 (0.25) N = 7079	4.87 (0.05) N = 7079	10.36 (0.12) N = 7079
Months 7-12	82.04 (0.24) N = 606	3.59 (0.16) N = 606	22.87 (1.05) N = 606	53.03 (0.31) N = 3961	3.91 (0.07) N = 3961	13.56 (0.25) N = 3961
Months 13-18	81.47 (0.35) N = 339	3.93 (0.23) N = 339	20.74 (1.24) N = 339	53.03 (0.41) N = 2192	4.39 (0.10) N = 2192	12.07 (0.29) N = 2192
Months 19-24	81.67 (0.37) N = 200	3.87 (0.29) N = 200	21.10 (1.61) N = 200	54.18 (0.58) N = 1181	4.39 (0.13) N = 1181	12.35 (0.39) N = 1181
Users initially enrolled with an annual contract						
Year 1	71.02 (0.50) N = 145	4.69 (0.38) N = 145	15.15 (1.24) N = 145	47.57 (0.75) N = 598	4.48 (0.17) N = 598	10.62 (0.44) N = 598
Year 2	73.78 (1.06) N = 36	6.85 (1.00) N = 36	10.77 (1.57) N = 36	50.09 (1.81) N = 112	6.59 (0.49) N = 112	7.60 (0.60) N = 112

Notes: Standard errors in parentheses. Standard errors for “Average price per average attendance” measure computed using the bivariate Delta method. The number of observations is denoted by N. An enrollment spell starts whenever an individual enrolls (or re-enrolls) in the club and ends whenever the individual quits or is censored. The sample “First spell” consists of the first enrollment spell. The sample “First spell and no subsidy” restricts the sample “First spell” to those spells in which the average adjusted monthly fee is at least \$70 if the spell starts with a monthly contract and at least \$58 if the spell starts with an annual contract.

For the 6-month periods, the sample includes spells that are ongoing, not frozen, and not miscoded in at least one month in the period. For year 1 in the annual contract, the sample includes only spells that started at least 14 months before the end of the sample period, and that were not prematurely terminated because of medical reasons or relocation. For year 2, the sample includes only spells that started with an annual contract at least 26 months before the end of the sample period, and that lasted at least 25 months. The spells in row “First contract monthly” start with a monthly contract. The spells in row “First contract annual” start with an annual contract. The “Average price” in period t is the average fee across people enrolled in period t. The “Average attendance” in period t is the average number of visits across people enrolled in period t. The measure in Column (3) is the ratio of the measure in Column (1) and the measure in Column (2).

Table 9: Loss from choice of flat-rate contracts

	Sample: First spell and no subsidy, first contract monthly, all clubs			Sample: First spell and no subsidy, first contract annual, all clubs		
	Join before October 1997 (1)	Join before April 1998 (2)	Join before October 1998 (3)	Join before October 1997 (4)	Join before April 1998 (5)	Join before October 1998 (6)
Loss from choice of flat-rate contract in \$	698.16 (106.05)	607.35 (53.33)	595.00 (41.40)	-61.65 (424.22)	220.17 (193.95)	230.01 (146.08)
Total money spent per spell in \$	1516.88 (114.96)	1309.07 (59.73)	1256.48 (46.15)	1832.34 (169.77)	1562.73 (101.25)	1445.00 (76.47)
Percentage of loss over money spent	47.87% (4.78)	51.28% (2.93)	51.83% (2.60)	6.01% (17.62)	28.45% (9.66)	27.18% (8.63)
Number of observations <i>N</i>	N = 70	N = 238	N = 345	N = 15	N = 43	N = 68

Notes: Standard errors in parentheses. The number of observations is denoted by *N*. An enrollment spell starts whenever an individual enrolls (or re-enrolls) in the club and ends whenever the individual quits or is censored. The sample “First spell” consists of the first enrollment spell. The sample “First spell and no subsidy” further restricts the sample “First spell” to those spells in which the average adjusted monthly fee is at least \$70 if the spell starts with a monthly contract and at least \$58 if the spell starts with an annual contract. The spells in Column “First contract monthly” start with a monthly contract. The spells in Column “First contract annual” start with an annual contract. The measure “Loss from choice of contract with flat fee in \$” is the average saving in \$ that a user who chose a contract would have attained if she had purchased a 10-visit pass for \$100 and attended the same number of times. A negative value denotes that the user would have lost money by purchasing the pass.