Online Appendix
Not for Publication

Switching Costs and Competition in Retirement Investment

Fernando Luco

A Additional Figures

Figure A.1: Distribution of fees over time

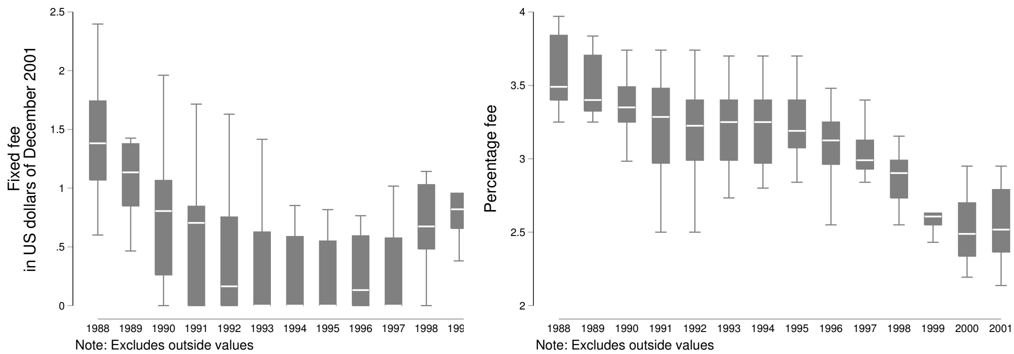

(a) Fixed fee (b) Percentage fee

The figures show the median (horizontal line), 25th, and 75th percentile (upper and lower extreme of the boxes), and the maximum and minimum adjacent values.

Figure A.2: Distribution of 36-months annualized realized returns and difference to industry mean

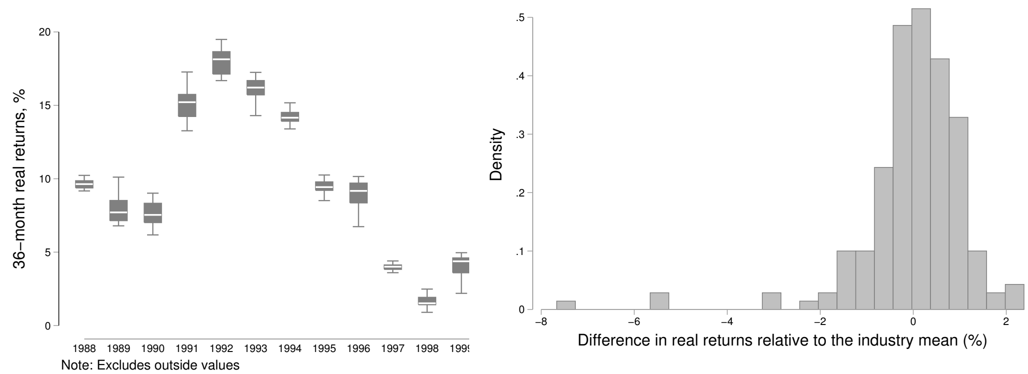

(a) Distribution of returns

(b) Returns relative to the industry mean

The figure on the left reports the median, 25th, 75th percentile, and the maximum and minimum adjacent values of the distribution of realized returns of each year. The figure on the right reports the distribution of returns relative to the industry mean for the entire period.

Figure A.3: Distribution of the number of times enrollees return to the system by year of enrollment

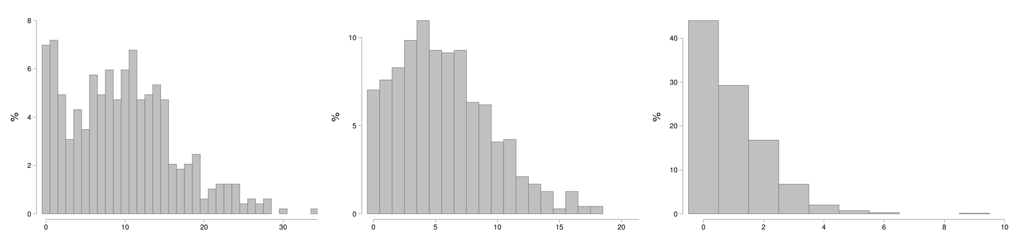

(a) 1988

(b) 1994

(c) 2000

B Additional Descriptive Evidence

To take advantage of the richness of the data, Table B.1 incorporates additional regressors to those in the probit regressions presented in Section II regarding which enrollees switch. Column 1 adds an indicator variable that takes the value of one for years 1998 to 2001. This follows because at the end of 1997, the regulator reformed the system and made it more difficult to switch in response to salespeople offering gifts to induce switching. As expected, such reforms had a negative effect on the probability of switching. Column 2 also controls for time elapsed before returning to the system and salary level. The results show that the longer a person was not participating before returning, the higher the probability of switching, though in this case the impact of changes in salary turns out to be negative and significant, while salary level is positively associated with the probability of switching. Column 3 replicates Column 2, dropping salary level and replacing it with account balance. It is shown that the probability of switching increases with account balance. Finally, Column 4 controls for all these variables jointly and shows that, not surprisingly, it is salary rather than the account balance's affecting the probability of switching.

TABLE B.1: More on the effect of demographics on the probability of switching

	(1)	(2)	(3)	(4)
Returning	0.783	0.687	0.703	0.688
	(0.0110)	(0.0125)	(0.0123)	(0.0125)
Increase in salary > 10%	-0.004	-0.0421	0.00194	-0.0407
	(0.009)	(0.0104)	(0.00981)	(0.0105)
Age	-0.009	-0.00989	-0.0100	-0.0100
	(0.001)	(0.000843)	(0.000849)	(0.000852)
Male	0.013	-0.00280	0.00452	-0.00333
	(0.013)	(0.0132)	(0.0133)	(0.0132)
Has a voluntary savings account	0.150	0.137	0.137	0.136
	(0.014)	(0.0138)	(0.0141)	(0.0139)
Year≥ 1998	-0.443			
	(0.013)			
Time elapsed before returning		0.0177	0.0173	0.0178
		(0.000947)	(0.000940)	(0.000948)
Salary (tens of thousands of Chilean pesos)		0.00440		0.00429
		(0.000319)		(0.000334)
Account balance (tens of thousands of Chilean pesos)			0.000441	0.0000406
			(0.0000696)	(0.0000733)
Marginal effect of returning (percentage points)	8.31	6.72	6.96	6.73
Year fixed effects	No	Yes	Yes	Yes
PFA fixed effects	Yes	Yes	Yes	Yes
N	341769	341769	341769	341769
Pseudo R^2	0.140	0.156	0.153	0.156
Log likelihood	-45272.9	-44410.6	-44545.1	-44410.3
χ^2	11554.0	12866.8	18755.7	12892.9

Note: Standard errors, clustered at the individual level, are in parentheses. The dependent variable is an indicator that is equal to one if the enrollee switches managers in that period and zero otherwise. A unit of observation is an enrollee in a month. Estimation corresponds to a probit regression. Marginal effects refer to the increase in the probability of switching, in percentage points, associated with the returning indicator taking the value of one relative to zero. All marginal effects are significant at the 1 percent level.

C Comparing Demographics Across Groups

The purpose of this section is to compare the demographic composition of the different groups enrollees may belong to: new, existing, and returning. Table C.1 shows that in terms of age and income, groups are remarkably similar. However, they do differ significantly on gender composition and whether enrollees have a voluntary savings account. Indeed, though 50% of the sample is female, males are overrepresented in the returning group. This means that after leaving the market, males are more likely to return to it than females, a fact that for years has worried the Chilean authorities because it results in low pensions for women. However, this is more closely is associated with the specifics of the labor market rather than selection in the pension-funds market. Furthermore, even if people were to return to the formal labor market because of a desire to participate in the retirement-investment system, this type of selection is not a concern from the perspective of identifying switching costs. Indeed, it would be necessary for people to transition between returning and existing because of relative preferences for a PFA–and not for the system–for selection to be a concern. Finally, existing enrollees are more likely to have a voluntary savings account. This suggests that there might be further differences across individuals that must be taken into account. I do so in estimation, as I control for observable individual characteristics, such as age, income, gender, and whether an enrollee has a voluntary savings account, and time-invariant individual-specific unobserved heterogeneity on a number of dimensions.

Table C.1: Demographics among returning and existing enrollees

	Year of initial enrollment					
	1988		1994		2000	
	Returning	Existing	Returning	Existing	Returning	Existing
Age	30	31.1	27.6	28.3	24.8	25.7
	(8.1)	(8.0)	(9.6)	(8.9)	(9.0)	(8.7)
Income	159947.7	167897.4	210069.2	192524.5	194350.9	203370.7
	(247714.5)	(190575.9)	(304510.3)	(177790.8)	(209558.2)	(198063.1)
Male	0.62	0.58	0.57	0.50	0.49	0.46
Voluntary savings account	0.31	0.36	0.22	0.26	0.07	0.08

Note: The table reports means and standard deviations for each variable depending on the year of enrollment. For each year, the average is taken across all observations that correspond to people either returning to the system or classified as existing enrollees.

D Additional results from the structural model

FIGURE D.1: Distribution of decision costs as a percentage of enrollment costs

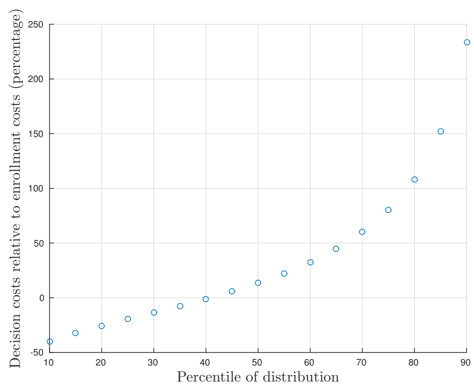

On the vertical axis, the figure reports decision costs as percentage of enrollment costs. A negative number means that enrollment costs are larger than decision costs. The horizontal axis plots percentiles of the distribution of the ratio of decision to mental costs. The figure shows that for 60 percent of enrollees, decision costs are larger than enrollment costs.

Table D.1: Robustness analysis (switching costs)

		(1)	(2)	(3)	(4)	(5)
Decision cost	Constant	3.488	3.621	2.913	3.449	3.43
		(0.069)	(0.069)	(0.449)	(0.069)	(0.070)
	Age	0.015	0.014	0.055	0.016	0.016
		(0.002)	(0.002)	(0.012)	(0.002)	(0.002)
	Male	-0.019	-0.021	0.026	-0.02	-0.019
		(0.034)	(0.034)	(0.007)	(0.034)	(0.034)
	Income	-0.004	-0.004	0.042	-0.004	-0.004
		(0.001)	(0.001)	(0.017)	(4E-4)	(5E-4)
	Voluntary savings	-0.383	-0.421	-0.312	-0.389	-0.393
		(0.036)	(0.036)	(0.050)	(0.036)	(0.036)
	Regulation	0.535	0.539	0.633	0.527	0.532
		(0.038)	(0.037)	(0.031)	(0.037)	(0.038)
	Balance	-0.0001				
		(2E-4)				
	Time unenrolled		-0.019			
			(0.002)			
	σ_{η_D}				0.043	0.043
					(0.011)	(0.011)
Enrollment cost	Constant	1.312	1.057	1.116	2.584	2.523
		(0.089)	(0.085)	(0.090)	(0.216)	(0.268)
	Age	0.009	0.014	0.051	0.018	0.017
		(0.003)	(0.003)	(0.010)	(0.003)	(0.003)
	Male	-0.009	0.003	0.058	0.022	0.021
		(0.041)	(0.041)	(0.009)	(0.045)	(0.045)
	Income	-0.009	-0.004	0.026	-0.010	-0.010
		(0.001)	(0.001)	(0.005)	(0.001)	(0.001)
	Voluntary savings	0.115	0.191	0.206	0.103	0.113
		(0.044)	(0.043)	(0.022)	(0.048)	(0.048)
	Regulation	0.345	0.403	0.397	0.710	0.702
		(0.048)	(0.047)	(0.039)	(0.065)	(0.074)
	Balance	0.001				
		(2E-4)				
	σ_{η_E}				2.004	1.943
					(0.157)	(0.209)
McFadden's Pseudo R^2		0.890	0.890	0.484	0.890	0.891
$-\left(\frac{1}{N}\mathcal{L}\right)$		0.267	0.267	1.295	0.266	0.266
Number of observations				350,660		

Note: Standard errors in parentheses. Column (1) includes balance in the specification of switching costs. Column (2) includes time unenrolled prior to returning in the specification of decision costs. Column (3) allows for autocorrelation in ε. Column (4) uses realized returns, rather than the ranking of returns, in the specification of demand. Finally, Column (5) follows a control function approach. PFA fixed effects included in all specifications. The dependent variable is an indicator that is equal to one for the chosen PFA and zero for the rest. Estimation is via maximum likelihood in columns 1 and 2, and simulated maximum likelihood in columns 4 and 5 (using 50 Halton draws per individual). Estimation in column 3 is done using frequencies after simulating draws of ε and computing the associated utility levels for each option in the choice set. In this case, standard errors are computed using 25 bootstrap samples drawn with replacement from the original dataset.

Table D.2: Robustness analysis (taste coefficients)

		(1)	(2)	(3)	(4)	(5)
Taste coefficient on income	Constant	3.9879	4.1265	4.1827	3.139	3.1632
		(0.3308)	(0.3297)	(0.1078)	(0.459)	(0.4629)
	Age	-0.0668	-0.0978	-0.0992	-0.022	-0.0230
		(0.0102)	(0.0097)	(0.0131)	(0.013)	(0.0137)
	Male	-0.58	-0.7367	-0.7356	-0.942	-0.9354
		(0.1615)	(0.1604)	(0.1583)	(0.201)	(0.2013)
	Balance	-0.0047				
		(0.0003)				
	σ_α				0.172	0.1661
					(0.069)	(0.068)
Taste coefficient on returns	Constant	-0.036	-0.045	-0.0464	-0.002	-0.0278
		(0.0060)	(0.0059)	(0.0044)	(0.0004)	(0.0068)
	Age	0.0009	0.0013	0.0018	0.00003	0.0009
		(0.0002)	(0.0002)	(0.0003)	(0.0009)	(0.0002)
	Male	-0.0064	-0.0041	-0.0037	0.0005	-0.0047
		(0.0031)	(0.0030)	(0.0035)	(0.0004)	(0.0034)
	Income	0.0003	0.0005	0.001	4E-5	0.0005
		(4.8E-005)	(4.6E-005)	(8.8E-005)	(3E-5)	(4.7E-5)
	Balance	0.0002				
		(1.3E-005)				
	σ_β				0.001	0.0045
					(0.003)	(0.0071)
$\lvert\rho\rvert$ or Control Function				0.0039		-3.7E-5
				(0.0884)		(9.8E-5)
McFadden's Pseudo R^2		0.890	0.890	0.484	0.890	0.891
$-\left(\frac{1}{N}\mathcal{L}\right)$		0.267	0.267	1.295	0.266	0.266
Number of observations				350,660		

Note: Standard errors in parentheses. Column (1) includes balance in the specification of switching costs. Column (2) includes time unenrolled prior to returning in the specification of decision costs. Column (3) allows for autocorrelation in ε. Column (4) uses realized returns, rather than the ranking of returns, in the specification of demand. Finally, Column (5) follows a control function approach. PFA fixed effects included in all specifications. The dependent variable is an indicator that is equal to one for the chosen PFA and zero for the rest. Estimation is via maximum likelihood in columns 1 and 2, and simulated maximum likelihood in columns 4 and 5 (using 50 Halton draws per individual). Estimation in column 3 is done using frequencies after simulating draws of ε and computing the associated utility levels for each option in the choice set. In this case, standard errors are computed using 25 bootstrap samples drawn with replacement from the original dataset.

E Additional material from the counterfactual exercises

E.1 Procedure used to simulate sequences of choices

For each of the counterfactuals described in Section V. A., I simulate the sequence of choices of a 30 percent random sample of enrollees. To do this, I simulate draws of ε's, compute the utility associated with the initial choice of each individual, then compute the sequence of choices for each and the resulting balances. To compute balances, I use the monthly contributions and price of a share of each fund for every month over the sample period. I limit the sample to 30 percent of the total number of enrollees, because the process is computationally demanding. Indeed, each counterfactual requires computing the sequence of choices for each individual over a large number of ε simulations and random draws for the taste coefficient (100 draws of ε for each individual and alternative, plus 50 Halton draws for each taste coefficient and switching cost). Increasing the sample size requires decreasing the number of draws that can be considered, does not provide additional insights, and does not affect the results.

E.2 Example of how switching costs affect choices

Figure E.1: Sequence of choices under different policy interventions

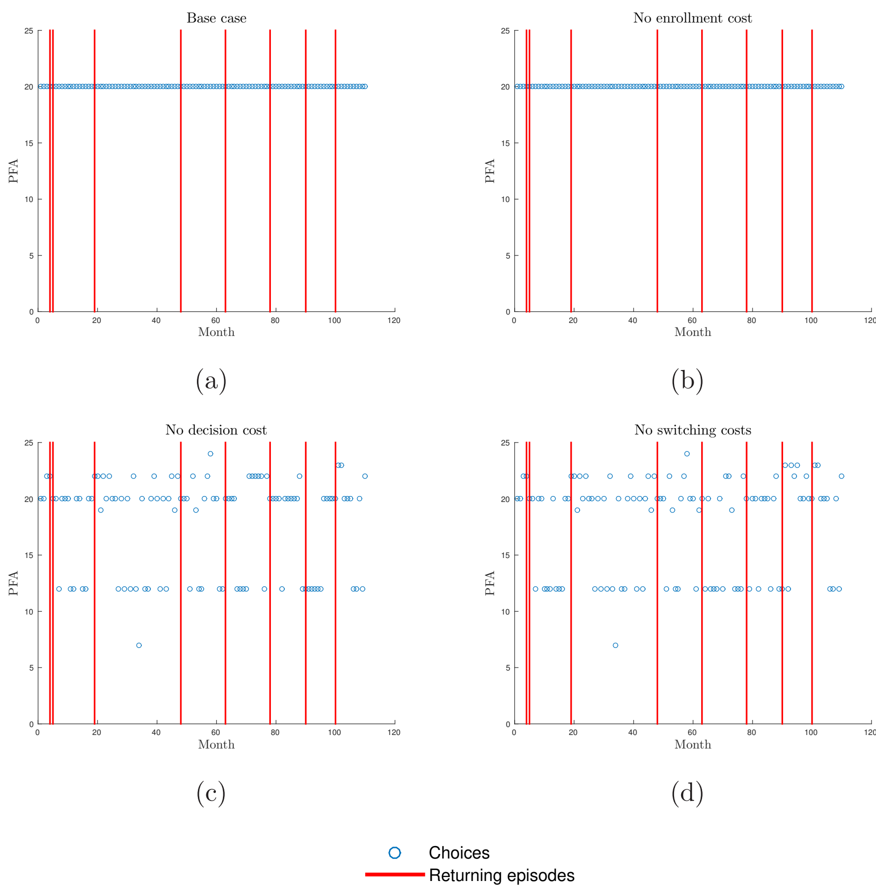

The figures report the enrollment history of a randomly chosen enrollee (random subject to the constraint of appearing at least 100 times). Each panel represents the enrollment history under a different counterfactual. Red vertical lines highlight periods in which the enrollee was a returning enrollee. In all other periods, with the exception of the first, the enrollee was an existing enrollee.

E.3 Dynamic Price Competition: Algorithm

This appendix describes the procedure followed when analyzing how switching costs affect equilibrium pricing. The starting point is Equation 5, which is reproduced below

$$(\star) \qquad \frac{\partial \Pi_{jt}}{\partial p_{jt}} + \beta \left[\frac{\partial \boldsymbol{s_t}}{\partial p_{jt}}\right]' \boldsymbol{E}_t \left[\frac{\partial V_j(\boldsymbol{s}_t, \boldsymbol{X}_t)}{\partial \boldsymbol{s}_t}\right] = 0.$$

Equation ($\star$) represents the set of first-order conditions that must to be solved, simultaneously, by the equilibrium fees. To solve the system of equation, it is necessary to compute the second term, and in particular, $\boldsymbol{E}_t\left[\frac{\partial V_j(\boldsymbol{s}_t, \boldsymbol{X}_t)}{\partial \boldsymbol{s}_t}\right]$, the expectation of the derivative of the value function with respect to the share vector. To compute this derivative, I follow a two-step procedure. First, I estimate a policy function $p = p(\boldsymbol{s}_t, \xi)$. With the policy function, I use the sequential representation of the value function and simulate N paths of length T. That is,

$$V(\boldsymbol{s}_t, \boldsymbol{X}_t) = \frac{1}{N} \sum_{i=1}^{N} \sum_{t=0}^{T} \beta^t \Pi(\boldsymbol{s}_t, \boldsymbol{p}_t, \boldsymbol{X}_t).$$

This approximation of the value function allows me to compute the continuation value for any given initial $\boldsymbol{s}_t$. Then, to compute the derivative of the value function, I use

$$(1) \qquad \frac{\partial V(\boldsymbol{s}_{t+1}, \boldsymbol{X}_t)}{\partial s_{jt}} = \frac{V(\boldsymbol{s}_t + \epsilon \boldsymbol{l}, \boldsymbol{X}_t) - V(\boldsymbol{s}_t - \epsilon \boldsymbol{l}, \boldsymbol{X}_t)}{2\epsilon},$$

where ϵ is a small constant and $\boldsymbol{l}$ is a vector (of the same length as $\boldsymbol{s}_t$) with a one in position l and zeros everywhere else. That is, the derivative of the value function with respect to share is computed by definition.

The approximation just described allows me to solve for equilibrium fees for any given starting vector of shares $\boldsymbol{s}_t$. Furthermore, once the space of shares has been discretized, the equation can be solved independently for all points in the grid. Finally, though the share grid may be coarse, this does not affect the computation of the derivative by using forward simulation, as the sequential representation of the value function results in shares, period to period, that need not be in the grid of points originally defined.

As described above, once the state space has been discretized, the problem can be solved independently for every point in the grid. However, the problem is still difficult from a computational perspective, as the derivative of the value function must be computed, by forward simulation, in all iterations during the search for optimal fees. For this reason, I make use of the Open Science Grid (Pordes et al., 2007; Sfiligoi et al., 2009) to compute equilibrium fees for each point in the grid of initial states (each initial point in the grid defines a different job submitted to the OSG). Then, I interpolate the resulting grid of equilibrium fees, to compute fees for a finer grid of states. Finally, to compute steady-state fees, I draw $10,000$ random initial states from the finer grid (the interpolated one), and for each randomly drawn state I search for the associated optimal fees within the array just described. With these fees, I recompute the implied shares and iterate until fees (and shares) do not change between two iterations. The fees reported in Subsection V. B. correspond to the mean expected fee computed across the 10,000 simulations using the implied shares to compute the expected fee for each simulation.

References

Pordes, Ruth, Don Petravick, Bill Kramer, Doug Olson, Miron Livny, Alain Roy, Paul Avery, Kent Blackburn, Torre Wenaus, W Frank, et al. 2007. "The open science grid." Vol. 78, 012057, IOP Publishing.

Sfiligoi, Igor, Daniel C Bradley, Burt Holzman, Parag Mhashilkar, Sanjay Padhi, and Frank Wurthwein. 2009. "The pilot way to grid resources using glideinWMS." Vol. 2, 428–432, IEEE.