Economics Bulletin

Volume 36, Issue 3

Loss Aversion and Student Achievement

David M McEvoy
Appalachian State University

Abstract

We conduct a field experiment to test if loss aversion behavior can be exploited to improve student performance in an undergraduate statistics course. In one treatment (gains), student grades were reported as points gained, and in the other treatment (losses) grades were reported as points lost. When controlling for other factors that affect student performance, we find that students in the loss treatment earned statistically higher grades than students in the gain treatment. Although preliminary, the results suggest that a simple manipulation of how grades are framed in the classroom can be a costless way to exploit loss aversion behavior and lead to higher student achievement.

Citation: David M McEvoy, (2016) "Loss Aversion and Student Achievement", *Economics Bulletin*, Volume 36, Issue 3, pages 1762-1770
Contact: David M McEvoy - mcevoydm@appstate.edu.
Submitted: January 12, 2016. **Published:** September 11, 2016.

1. Introduction

Education opens doors. The more successful students are in school, the more options they will have later on in life. Student performance, most often measured by objective test scores, depends on a number of factors, one of which is the level of effort students put forth to learn the material. Motivating students to work harder in school, therefore, is a worthwhile pursuit. A large literature on loss aversion and reference dependent preferences suggests that people are more responsive to incentives framed as losses than to incentives framed as gains (Kahneman and Tversky 1979; Tversky and Kahneman 1991; Ariely et al. 2005). We conduct a field experiment to test if loss aversion behavior can be exploited to improve student performance in an undergraduate statistics course. To do so, we manipulate the framing of the grading scales across statistics classes in the same semester. In one treatment, the class starts with zero points and students are given the opportunity to gain up to 560 points. In the other, the class starts with 560 points (100 percent) and students can only incur losses from their performance. Our hypothesis, informed by theories of loss aversion and reference dependent preferences (e.g., Köszegi and Rabin 2006), is that students that start with 100 percent will exert more effort to keep their high grades in comparison to students that start with zero and work their way up from low grades. Although we only imperfectly observe effort levels (i.e., attendance), we observe outcomes (grades) that depend on effort levels.

Related behavioral research has explored the use of material incentives (money and trophies) to exploit loss aversion related to student performance (Fryer et al. 2012; Levitt et al. 2012).[1] In these studies rewards are either paid in response to good student performance (gains) or taken away in response to poor performance (losses). The results from Fryer et al. (2012), in which teachers earned/lost rewards conditional on student performance, suggest that the incentive structure matters; students performed significantly better when teachers were under the threat of losing a financial reward compared to the promise of earning an equivalent one. On the other hand, Levitt et al. (2012) find no difference in performance when students could either gain or lose material rewards. Although our study is closely related to this literature, our experiment does not rely on material incentives. The experimental design simply varies the frame in which students' grades are reported throughout a semester. The approach, therefore, has the potential to be a very cost-effective method to incentivize students. Our motivating hypothesis is based on the idea that preferences, in part, are reference dependent. Our theory presumes that students have a prior subjective range of "acceptable grades" in the course, and that when students are in this range they will work relatively harder to avoid a lower grade than they will to earn a higher one.

Our overarching hypothesis and experimental design is similar to a recently published study by Apostolova-Mihaylova et al. (2015). Their study, like ours, was conducted during the 2012-2013 academic year.[2] In contrast to our study, they did not find a significant treatment effect between their treatment group (losing points) and the control group (earning points). However, they do find a significant difference in how males and females reacted to the experimental frame. They show that men in the treatment group scored significantly higher than

[1] A much larger literature explores the effects of material and non-material rewards on student performance without appealing to the behavioral economics of loss aversion (e.g., Angrist and Lavy 2009; Azmat and Iriberri 2010; Bettinger 2010; Leuven et al. 2010; Fryer 2011, 2013)

[2] Our study and Apostovola-Mihaylova et al. (2015) were conducted concurrently, but without knowledge of each other's experiment or research agenda.

men in the control group. Conversely, women earned significantly lower grades in the treatment group.

Our results differ in important ways from Apostolova-Mihaylova et al. (2015). Consistent with our hypothesis, we find that students part of the "loss" treatment (starting with 560 points) earn statistically higher grades than students in the "gain" treatment (starting with zero points). Also consistent with a theory of reference dependent preferences, we observe changes in the relative performance between the two treatments over time. Most importantly, students in the loss treatment were more likely to attempt the final exam and their average grade on the final was statistically higher. We do not observe significant differences across gender.

2. Theory and experimental design

Our hypotheses are derived from the established theory on loss aversion and reference dependent preferences from psychology and behavioral economics (Tversky and Kahneman 1991). We, therefore, do not offer a formal theory. However, because participants in our experiment start off in the extremes of possible outcomes, it is useful to characterize how reference points play a role in relative effort levels and, indirectly, student performance.

Figure 1 shows a grading scale from zero to 100 percent. Students either work their way up from zero or down from 100. The rectangular bars represent the effort levels exerted by students within the different grading ranges. Students starting at zero, exert a relatively high level of effort to move from failure into their subjective range of acceptable grades (from some lower threshold to 100 percent). Once a student that starts from below crosses the threshold into her acceptable range her effort level drops. That is, once in the acceptable range she exerts less effort to earn a higher acceptable grade than she would to move from an unacceptable to acceptable one. Students that start with 100 percent, on the other hand, exert effort to avoid losing their high grade. Leaning on the theory of loss aversion, students within the acceptable grade range are expected to exert less effort to make gains in their grade relative to the corresponding effort level students put forth to avoid losses.

Figure 1: Loss aversion and student effort

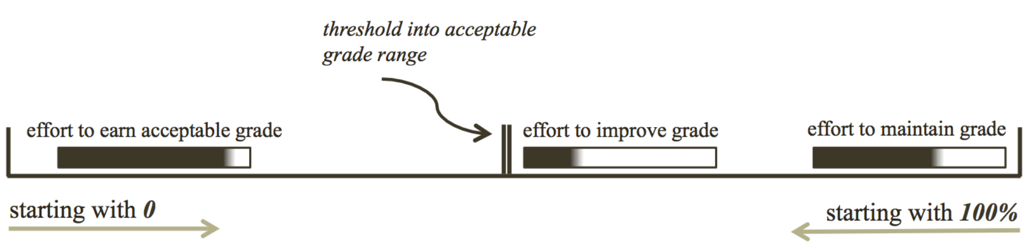

The experiment was conducted during the Fall 2012 academic semester at Appalachian State University. The course number (ECO 2200), instructor, materials and assessment activities were held constant across the three sections, only the framing of the grading scale varied. The assessment activities included quizzes (seven), problem sets (four), exams (four) and active participation (iClicker performance). In one class - the gain treatment - students started with zero points and earned points up to 560, and in the other two classes - the loss treatment - students started with 560 points and lost points conditional on performance. A student's percentage grade was calculated as 100*(Total Points/560). The grading interface was constructed in Excel and an active link to the grade book was embedded into the course website. With both grading formats, the students' current number of points (lost or gained), percentage grade and letter grades was continuously updated and posted throughout the semester.

The 560 points were broken down as follows. There were three in-semester exams and a comprehensive final exam that could be used to replace the lowest of the three in-semester exams. The three highest exams accounted for 300 points (each about 18 percent of the course grade). There were seven quizzes, each worth 20 points, and the lowest quiz grade was dropped. Quizzes therefore accounted for 120 points. Four problem sets each accounted for 20 points (80 in total) and active participation accounted for 60 points. Note that active participation was made up of 1/3 attendance and 2/3 performance using classroom clickers.

One element of the grading scheme that is particularly important is how the final exam plays its role. The final exam is offered during a week after the semester comes to a close, so when a student debates whether to take the final exam all other grades have been computed. Therefore, students in the gain treatment are shown their semester grade given a final exam score of zero. On the other hand, students in the loss treatment are shown their grade given a final exam score of 100. Students, in either treatment, could plug in their expected grade for the final exam to view their expected course grade.

Our sample consisted of 177 business students that each attended one of three statistics courses. The duration of the courses was one hour and fifteen minutes and they were held back-to-back-to back, with start times of 9:30, 11:00 and 12:30. Students were free to register for any of the three courses once their registration period opened.[3] The 9:30 and 12:30 classes served as the loss treatments (starting with 560 points), and the 11:00 class served as the gain treatment (starting with 0 points). There were 56, 66, and 55 students in the 9:30, 11:00 and 12:30 classes, respectively. The same classroom was used for all three courses and spaces remained open in all three throughout the semester (registration was capped at 72). In total, 111 students participated in the loss treatment and 66 students participated in the gain treatment. Overall, sixty eight percent (121) of the students were Male.[4]

3. Results

We want to test whether students that started with 100 percent and received their grading feedback framed as losses performed better than students starting with zero and working their way up. When analyzing the unconditional results using a pairwise t-test to compare overall course grades, we do not find a significant effect across treatments. The average grades were an 81.7 and an 82.19 in the loss and gain treatments respectively (p = 0.696).[5,6] The study was

[3] The opening of the registration period was scattered for students depending on how many credit hours they have earned. Students with more credits had priority in the sense that their registration period started earlier. All three courses had a similar pattern of enrollment throughout the entire registration period and none of the classes reached the enrollment cap of 72.

[4] We recognize that the study would have benefited from more balanced sample sizes between the two treatments. Given the teaching assignment constraints this was not possible.

[5] Apostolova-Mihaylova et al. (2015) report no significant difference in average grades between their control (grades framed as gains relative to zero) and their treatment (grades framed as losses relative to 100) when not

designed, however, to control for other observable factors that could affect performance independent of how the grading scheme was framed. We collected data on students' GPA, SAT score, attendance, their numeric grade from their previous statistics course, the number of hours they study per week (for all classes combined), gender and whether they attempted the final exam.[7] We did not include a variable for their particular major because our sample is limited to business or aspiring business majors and many of them only report intended majors at this stage in their academic career.

Table I reports the results from linear models regressing a student's percentage grade for the course on the aforementioned variables. From the initial 177 observations, 32 were dropped due to missing observations in one or more of the explanatory variables. Of the remaining 145 observations, 56 were from the gain treatment and 89 were from the loss treatment (47 and 42 students in the early and late section, respectively).[8] The first column of results in Table I is from a pooled model (all 145 observations) regressing course grade on the treatment dummy and the control variables, excluding the dummy variable whether the student attempted the final. Recall that the nature of the final exam in this class is particularly important. Students in the gain treatment decided whether to attempt the final knowing it can only improve upon their current course grade (replacing lowest grade), whereas students in the loss treatment had a default final course grade of 100 before attempting the final, and therefore earning a grade lower than 100 on the final would likely lower their current course grade. We run separate models to better isolate the effect attempting the final exam has on course grades.

---

controlling for the interaction between gender and treatment. However, their result is not directly comparable because they do not report the results from a two-sample test on the unconditional means.

[6] The Wilcoxon rank-sum test also indicates that the two samples come from identical populations ($p = 0.893$).

[7] GPA, SAT, previous statistics grade and hours studying were self-reported through a survey administered through the course website. The number of classes attended was gathered from the iClicker system.

[8] A t-test comparing the unconditional average grades with the 145 remaining observations again reveals that there is no statistical difference between the gain and loss treatment ($p = 0.650$). Moreover, there is no statistical difference in average grades between the two classes (early and late) that make up the loss treatment ($p = 0.355$).

Table I: Course grades as a function of treatment and control variables

	Pooled	Pooled	Gains	Losses
Loss Treatment	1.961*	2.492**	---	---
	(1.109)	(1.103)		
Attendance	0.632***	0.607***	0.407	0.698***
	(0.205)	(0.200)	(0.427)	(0.237)
GPA	8.089***	6.999***	5.828*	7.726***
	(1.593)	(1.611)	(2.994)	(1.960)
Stats 1 Grade	0.250***	0.254***	0.343***	0.225**
	(0.083)	(0.081)	(0.132)	(0.106)
SAT	-0.0002	-0.000	-0.006	0.000
	(0.001)	(0.001)	(0.004)	(0.001)
Hours Studying	0.136	0.096	-0.094	0.210*
	(0.086)	(0.086)	(0.148)	(0.108)
Female	-0.964	-0.803	0.451	-1.346
	(1.201)	(1.177)	(1.941)	(1.541)
Attempted Final	---	-3.036***	-3.634*	-2.138
		(1.137)	(1.892)	(1.513)
Constant	21.857***	27.005***	27.934***	26.523***
	(7.461)	(7.549)	(10.079)	(9.711)
N	145	145	56	89
R-squared	0.414	0.443	0.452	0.472
F	13.80	13.50	5.65	10.33

Note: Standard errors are in parentheses and *, **, *** indicate significance at the 0.10, 0.05 and 0.01 levels respectively.

From the first column in Table I we see that conditioning for other potential factors that affect students' performance, framing grades as losses as opposed to gains leads to statistically higher average course grades. The effect is weak, however, only significant at the 0.10 level. We also see that attendance, GPA, and Stats 1 grade have a significant and positive effect on performance at the 0.01 level.[9]

When we include the dummy variable for whether the student attempted the final exam (second column in Table I), the qualitative impact of the loss frame remains but with greater significance (now at the 5% level). The overall result suggests that a simple, and costless, twist to a grading scheme can successfully exploit loss-aversion behavior and lead to improved student performance. The estimate from the second model in Table I suggests that students' grades are on average 2.492 percent points higher in the loss treatment compared to the gain treatment, and this difference is significant ($p = 0.031$).

We also see that attempting the final exam is associated with significantly lower course grades. To explore this result further we estimate regression models stratified by treatment. From the third and fourth columns in Table I we observe that attempting the final is associated with lower course grades in the gain treatment (at the 0.10 level) but is not associated with a significant change in course grades in the loss treatment. Perhaps this is because in the gain

[9] To rule out the possibility of a section effect between the two classes that combine to form the loss treatment, we ran an additional regression using only data from the loss treatment. We regressed course grades on the remaining explanatory variables along with a section dummy variable. The coefficient was insignificant ($p = 0.316$) indicating the two sections are statistically equivalent.

treatment those attempting the final are more likely to be struggling students (those trying to reach their acceptable grade range) whereas in the loss treatment there is more heterogeneity in the types of students attempting the exam. We will explore the final exam further in a moment.

As illustrated in Figure 1, we hypothesize that students are more responsive to potential losses than they are to gains once they reach their acceptable grade range. Given this conjecture, we should expect to see more dramatic differences in effort (and resulting grades) at the end of the semester when students that started from zero have potentially crossed the threshold into their acceptable grade range. The column headings in Table II contain all of the graded assignments given throughout the semester (*PS* denotes problem sets, *Q* denotes quizzes and *E* denotes exams). The table includes the average course grades for each assignment by treatment. For each assignment we estimated a linear regression using the same treatment and control variables as the model presenting in the first column of Table I. Since there are 15 regressions, we chose to simplify the exposition and only report on the significance of the dummy variable for the loss treatment. The only assignments in which the average grades were statistically higher in the loss treatment were Problem Set 3 (second to last problem set), Quiz 7 (last quiz), Exam 3 (last day of class) and the Final Exam (finals week). Significance levels are denoted with asterisks.

Table II: Average grades by assignment and treatment

	PS1	*PS2*	*PS3**	*PS4*	*Q1*	*Q2*	*Q3*	*Q4*	*Q5*	*Q6*	*Q7***	*E1*	*E2*	*E3****	*Final**
Gains	18.39	16.70	15.39	18.24	16.03	16.39	15.70	15.67	14.79	17.03	17.76	72.39	72.91	75.33	63.74
Losses	18.23	16.25	16.16	18.09	16.31	15.19	14.22	15.62	15.04	17.55	18.58	69.62	74.38	78.34	68.71

Notes: Students' grades for each individual assignment listed in Table II were regressed (using OLS) on the same treatment and control variables from the first model in Table I. The symbols *,**,*** indicate that the grades in the loss treatment were significantly higher than grades in the gain treatment at the 10, 5 and 1 percent levels respectively. The grades for all other assignments were not statistically different between the two treatments. The coefficient estimates were suppressed for a cleaner exposition of the main findings.

We again turn our attention to the final exam. Recall that before students attempt the final they have full information regarding all of the other graded assignments. For those in the gain treatment, their course grades are posted given a final exam score of zero, while those in the loss treatment see their course grades given a final exam score of 100 percent.

In total, 73 (66 percent) and 31 (47 percent) students attempted the final in the loss and gain treatments, respectively. The two percentages are statistically different at the 0.01 level indicating that a larger fraction of students attempted the final when their grades were framed as losses as opposed to gains. To check the robustness of this result, we estimate a series of probit models where the dependent variable takes on a value of one if the student attempted the final (and a zero otherwise). We use the same treatment and controls as the regressions in Table I with the addition of a variable containing the student's total points before attempting the final. The first model uses the pooled dataset (145 observations) and the results are contained in Table III. We find that students are more likely to attempt the final in the loss treatment (significant at the 0.01 level). We also see that in the pooled model, the more points a student has earned the less likely they are to attempt the final. Again, we stratify the sample by the two treatments and see that this effect is qualitatively the same for both treatments. However, it is less significant in the loss treatment.

Table III: Probit regressions on the likelihood of attempting the final exam

	Pooled	Gains	Losses
Loss Treatment	1.959***	---	---
	(0.496)		
Attendance	0.006	0.302**	-0.100
	(0.054)	(0.136)	(0.075)
GPA	-0.600	-1.209*	-0.471
	(0.379)	(0.681)	(0.493)
Stats 1 Grade	0.020	0.012	0.022
	(0.019)	(0.034)	(0.025)
SAT	0.0001	-0.001	0.0002
	(0.0002)	(0.001)	(0.0003)
Hours Studying	-0.033	-0.004	-0.042
	(0.021)	(0.042)	(0.026)
Female	0.075	-0.544	0.266
	(0.272)	(0.478)	(0.378)
Points before final	-0.014***	-0.018***	-0.014**
	(0.004)	(0.007)	(0.007)
Constant	5.674***	5.768*	9.069***
	(1.952)	(3.319)	(3.446)
N	145	56	89
chi-squared	42.65	20.88	22.90
p-value (chi-squared)	0.000	0.004	0.002

Note: Standard errors are in parentheses and *, **, *** indicate significance at the 0.10, 0.05 and 0.01 levels respectively.

Finally, we regress the final exam scores on the same treatment and control variables from Tables I and II. In the first model we include the variables from the initial set of regressions in Table I. Here (column one, Table IV) we observe a positive effect on final exam grades from the loss treatment, but only at the 0.10 level. When we include the variable for total points earned before the final (column two, Table IV), the treatment effect is no longer significant. It appears that the loss-treatment effect on final exam scores is weak, and highly sensitive to how well students are doing prior to taking the exam. When stratifying by treatment, we see that the more points students have before the final exam the better they do on the final, but only in the loss treatment.

Table IV: Final exam grade as a function of treatment and control variables

	Pooled	Pooled	Gains	Losses
Loss Treatment	**5.046***	-5.857	---	---
	(2.823)	(4.752)		
Attendance	-0.257	-0.644	-3.074	-0.769
	(0.429)	(0.433)	(1.282)	(0.503)
GPA	5.381	2.443	11.479	-0.953
	(3.878)	(3.858)	(7.837)	(5.049)
Stats 1 Grade	0.234	-0.005	0.141	-0.030
	(0.213)	(0.221)	(0.522)	(0.257)
SAT	-0.0002	-0.0004	-0.003	-0.0006
	(0.002)	(0.001)	(0.012)	(0.002)
Hours Studying	0.304	0.200	0.125	0.228
	(0.252)	(0.244)	(0.509)	(0.308)
Female	1.727	3.018	6.009	2.492
	(2.877)	(2.792)	(6.413)	(3.448)
Points before final	---	**0.097*****	0.079	**0.109****
		(0.035)	(0.064)	**(0.049)**
Constant	29.238*	30.332*	7.251	33.955
	(18.022)	(17.254)	(41.738)	(21.178)
N	82	82	23	59
R-squared	0.125	0.209	0.393	0.157
F	1.50	2.41	1.39	1.36

Note: Standard errors are in parentheses and *, **, *** indicate significance at the 0.10, 0.05 and 0.01 levels respectively.

4. Conclusion

We manipulated the grading scheme across multiple statistics classes in the same semester to test if loss-aversion behavior can be exploited to improve student performance. Our design consisted of two treatments, one in which students started with 100 percent of total points and received feedback in “points lost” and another in which students started with zero and received feedback in “points earned”. Consistent with our hypothesis we find that when controlling for other factors that may affect student performance, students that were part of the loss treatment earned statistically higher grades than students in the gain treatment. While robust to different model specifications, the difference is only weakly significant (ranging from 0.05 to 0.10 levels).

Apostolova-Mihaylova et al. (2015) was the first published study to examine how framing grades as gains compared to losses can affect student performance. Our study was conducted during the same academic year (without knowledge of each others’ projects) and our results differ in meaningful ways. First, we do not observe a gender effect. Second, we observe higher average grades when grades are framed as losses compared to gains when controlling for other factors that potentially affect performance. Third, we observe a timing effect in which the impact of framing grades as losses is more influential later in the semester. Fourth, students in the loss treatment were significantly more likely to attempt the final exam, which replaces the lowest of the three semester exams. Although preliminary, our results suggest that a simple manipulation of how grades are framed in the classroom can be a costless way to exploit loss aversion behavior and lead to greater student achievement.

References

Angrist, J.D and V. Lavy (2009) "The Effect of High-Stakes High School Achievement Awards: Evidence from a Randomized Trial." *American Economic Review* 99(4): 1284-1414.

Apostolova-Mihaylova, M., W. Cooper, G. Hoyt and E.C. Marshall (2015) "Heterogeneous Gender Effects under Loss Aversion in the Economics Classroom: A Field Experiment." *Southern Economic Journal* 81(4): 980-94.

Ariely, D., J. Huber and K. Wertenbroch (2005) "When Do Losses Loom Larger than Gains?" *Journal of Marketing Research* 42(2): 134-38.

Azmat, G. and N. Iriberri (2010) "The Importance of Relative Performance Feedback Information: Evidence from a Natural Experiment Using High School Students." *Journal of Public Economics* 94(7-8): 435-452.

Bettinger, E. (2010) "Paying to Learn: The Effect of Financial Incentives of Elementary School Test Scores." *NBER Working Paper Series*, Working paper 16333.

Fryer, R.G. (2011) "Financial Incentives and Student Achievement: Evidence from Randomized Trials." *The Quarterly Journal of Economics* 126(4): 1755-1798.

Fryer, R.G. (2013) "Teacher Incentives and Student Achievement: Evidence from New York City Public Schools." *Journal of Labor Economics* 31(2): 373-427.

Fryer, R.G, S.D. Levitt, J. List and S. Sadoff (2012) "Enhancing the Efficacy of Teacher Incentives through Loss Aversion: A Field Experiment." *NBER Working Paper Series*, Working paper 18237.

Kahneman, D. and A. Tversky (1979) "Prospect Theory: An Analysis of Decision Under Risk." *Econometrica* 47(2): 263-92.

Kőszegi, B. and M. Rabin (2006) "A Model of Reference-dependent Preferences." *The Quarterly Journal of Economics* 121(4): 1133-1165.

Levitt, S.D., J.List, S. Neckermann and S. Sadoff (2012) "The Behavioralist Goes to School: Leveraging Behavioral Economics to Improve Educational Performance." *NBER Working Paper Series*, Working paper 18165.

Leuven, E., H. Oosterbeek and B.van der Klaauw (2010) "The Effect of Financial Rewards on Students' Achievements: Evidence from a Randomized Experiment." *Journal of the European Economic Association* 8(6): 1243-1265.

Tversky, A. and D. Kahneman (1991) "Loss Aversion in Riskless Choice: A Reference-Dependent Model." *The Quarterly Journal of Economics* 106(4): 1039-61.