Using Goals to Motivate College Students: Theory and Evidence from Field Experiments

Web Appendix

Damon Clark, Department of Economics, UC Irvine and NBER
David Gill, Department of Economics, Purdue University
Victoria Prowse, Department of Economics, Purdue University
Mark Rush, Department of Economics, University of Florida

Web Appendix I Tables

	All semesters	Fall 2013 & Spring 2014 (Performance-based goals)	Fall 2014 & Spring 2015 (Task-based goals)
Number of participating students	3,971	1,967	2,004
Number of students in Treatment group	1,979	995	984
Number of students in Control group	1,992	972	1,020
Fraction of students in Treatment group	0.50	0.51	0.49

Notes: The number of participating students excludes: students who did not give consent to participate; students who formally withdrew from the course; students who were under eighteen at the beginning of the semester; students for whom the university's Registrar data does not include SAT or equivalent aptitude test scores; and one student for whom the Registrar data does not include information on gender.

Table A.1: Participant numbers and treatment rates.

	Mean value		Treatment-Control difference		
	Treatment group	Control group	Difference	S.E.	p-value
Age	0.005	-0.005	0.010	0.032	0.764
Male	0.477	0.455	0.022	0.016	0.171
Black	0.064	0.051	0.012	0.007	0.091
Non-Hispanic white	0.604	0.619	-0.015	0.015	0.328
Hispanic	0.193	0.192	0.000	0.013	0.984
Asian	0.102	0.092	0.009	0.009	0.328
SAT score	0.001	-0.001	0.002	0.032	0.945
High school GPA	-0.016	0.016	-0.032	0.032	0.320
Advanced placement credit	0.759	0.756	0.003	0.014	0.800
Fall semester	0.620	0.607	0.012	0.015	0.435
First login time	-0.004	0.004	-0.007	0.032	0.820

Notes: The Treatment and Control groups contain 1,979 and 1,992 students respectively. Information about age, gender, race, SAT scores, high school GPA and advanced placement credit was obtained from the university's Registrar data. Age is measured on the first of the month in which the semester began and is rounded down to the nearest whole month. The variable SAT score is the sum of the student's scores on the verbal, analytic and numerical components of the primary aptitude test in the Registrar data (these are SAT scores for the majority of students). Advanced placement credit is an indicator for the student having entered the university with advanced placement credit. Fall semester is an indicator for the student having participated in the course in the Fall semester. First login time is the elapsed time between when the first email invitation to take the syllabus quiz was sent and when the student first logged into the course webpage. Each of the non-binary characteristics (age, SAT score, high School GPA and first login time) has been standardized to have a mean of zero and a variance of one within the Fall 2013 and Spring 2014 semesters combined (the performance-based goals experiment) and within the Fall 2014 and Spring 2015 semesters combined (the task-based goals experiment). The standardization of SAT score is stratified to ensure that this variable has the same mean and the same variance among students taking each type of aptitude test. S.E. is the standard error of the difference between the characteristic mean in the Treatment group and the characteristic mean in the Control group and is obtained assuming independent samples with equal variances. p-value is the two-sided p-value for the null hypothesis that the magnitude of the difference between the characteristic mean in the Treatment group and the characteristic mean in the Control group is zero. The joint significance of the characteristics is tested using a χ-squared test based on the results of a probit regression of an indicator for treatment on an intercept and the eleven characteristics listed in this table: the p-value for the joint null hypothesis that none of the eleven characteristics predicts treatment is 0.636.

Table A.2: Characteristics of students across all semesters

	Mean value		Treatment-Control difference		
	Treatment group	Control group	Difference	S.E.	p-value
Age	0.007	-0.007	0.014	0.045	0.761
Male	0.491	0.457	0.035	0.023	0.124
Black	0.069	0.047	0.022	0.011	0.037
Non-Hispanic white	0.628	0.644	-0.016	0.022	0.464
Hispanic	0.175	0.175	0.000	0.017	0.999
Asian	0.094	0.085	0.009	0.013	0.482
SAT score	-0.048	0.050	-0.098	0.045	0.030
High school GPA	-0.044	0.045	-0.089	0.045	0.049
Advanced placement credit	0.762	0.770	-0.008	0.019	0.686
Fall semester	0.575	0.591	-0.016	0.022	0.482
First login time	-0.003	0.004	-0.007	0.045	0.876

Notes: The Treatment and Control groups contain 995 and 972 students respectively. The p-value for the joint null hypothesis that none of the eleven characteristics predicts treatment is 0.153. Also see the notes to Table A.2.

Table A.3: Characteristics of students in Fall 2013 & Spring 2014 semesters (performance-based goals experiment)

	Mean value		Treatment-Control difference		
	Treatment group	Control group	Difference	S.E.	p-value
Age	0.003	-0.003	0.005	0.045	0.903
Male	0.462	0.454	0.008	0.022	0.704
Black	0.058	0.055	0.003	0.010	0.769
Non-Hispanic white	0.579	0.595	-0.016	0.022	0.472
Hispanic	0.210	0.209	0.002	0.018	0.932
Asian	0.109	0.099	0.010	0.014	0.476
SAT score	0.051	-0.049	0.101	0.045	0.024
High school GPA	0.013	-0.012	0.025	0.045	0.579
Advanced placement credit	0.756	0.742	0.014	0.019	0.472
Fall semester	0.665	0.624	0.041	0.021	0.055
First login time	-0.004	0.004	-0.007	0.045	0.868

Notes: The Treatment and Control groups contain 984 and 1,020 students respectively. The p-value for the joint null hypothesis that none of the eleven characteristics predicts treatment is 0.471. Also see the notes to Table A.2.

Table A.4: Characteristics of students in Fall 2014 & Spring 2015 semesters (task-based goals experiment)

	All students in the task-based goals experiment		All students in the performance-based goals experiment	
	Total points score		Total points score	
	OLS	Median	OLS	Median
Effect of asking students to set task-based goals	0.743 (0.474) [0.117]	0.924* (0.475) [0.052]		
Effect of asking students to set performance-based goals			-0.237 (0.458) [0.605]	-0.360 (0.494) [0.466]
Effect / (SD in Control group)	0.068	0.085	-0.022	-0.034
Mean of dependent variable in Control group	83.111	83.111	83.220	83.220
Observations	2,004	2,004	1,967	1,967

Notes: The regressions are the same as those reported in Table 4, except that we no longer include controls for student characteristics. Heteroskedasticity-consistent standard errors are shown in round brackets and two-sided p-values are shown in square brackets. *, ** and *** denote, respectively, significance at the 10%, 5% and 1% levels (two-sided tests).

Table A.5: Effects of task-based goals and performance-based goals on student performance without controls for student characteristics

	Effect of task-based goals on the average level of task completion	Effect of task-based goals on student performance	Effect of performance-based goals on student performance
	Panel I: Original specification (Tables 2 and 4)		
Treatment effect	0.491**	0.742*	0.300
	(0.205)	(0.431)	(0.398)
	[0.017]	[0.086]	[0.452]
Effect/ (SD in Control Group)	0.102	0.068	0.028
	Panel II: Specification with treatment interacted with SAT score bins		
Average treatment effect	0.516**	0.752*	0.278
	(0.207)	(0.431)	(0.395)
	[0.012]	[0.081]	[0.481]
Average effect/ (SD in Control Group)	0.108	0.069	0.026
Treatment × SAT score bin 1	-0.283	1.420	0.402
	(0.430)	(1.096)	(0.985)
	[0.510]	[0.195]	[0.683]
Treatment × SAT score bin 2	0.933**	0.726	0.868
	(0.471)	(0.962)	(0.999)
	[0.048]	[0.451]	[0.385]
Treatment × SAT score bin 3	1.046**	0.715	0.497
	(0.469)	(0.997)	(0.880)
	[0.026]	[0.473]	[0.572]
Treatment × SAT score bin 4	-0.098	0.832	0.180
	(0.462)	(0.936)	(0.809)
	[0.832]	[0.374]	[0.824]
Treatment × SAT score bin 5	0.984**	0.067	-0.554
	(0.462)	(0.769)	(0.844)
	[0.033]	[0.931]	[0.512]
Mean of dependent variable in Control group	8.627	83.111	83.220
Observations	2,004	2,004	1,967

Notes: Panel I repeats the OLS regression results with controls for student characteristics from Tables 2 and 4. The first column of Panel II shows results from an OLS regression of the number of practice exams completed on an indicator for the student having been randomly allocated to the Treatment group in the task-based goals experiment interacted with indicators for SAT score bin (bin 1 contains the lowest twenty per cent of SAT scores, bin 2 contains SAT scores greater than the 20th percentile and less than or equal to the 40th percentile, and so forth). The second column of Panel II is the same as the first column except that the dependent variable is total points score. The third column of Panel II shows results from an OLS regression of total points score on an indicator for the student having been randomly allocated to the Treatment group in the performance-based goals experiment interacted with indicators for SAT score bin. All three regressions in Panel II include controls for SAT score bin, along with all of the controls described in the notes to Table 2 except for those controls that are based on SAT score. For each regression in Panel II, the average treatment effect is the average of the five coefficients on the treatment-SAT-score-bin interactions. Heteroskedasticity-consistent standard errors are shown in round brackets and two-sided p-values are shown in square brackets. Standard errors for the average treatment effect are calculated using the delta method. *, ** and *** denote, respectively, significance at the 10%, 5% and 1% levels (two-sided tests).

Table A.6: Effects of task-based goals and performance-based goals by SAT score bin

	Male students in the task-based goals experiment		Male students in the performance-based goals experiment	
	Total points score		Total points score	
	OLS	Median	OLS	Median
Effect of asking students to set task-based goals	1.581*	1.700**		
	(0.706)	(0.674)		
	[0.050]	[0.024]		
Effect of asking students to set performance-based goals			-0.223	0.041
			(0.672)	(0.665)
			[1.000]	[1.000]
Effect / (SD in Control group)	0.141	0.151	-0.021	0.004
Mean of dependent variable in Control group	83.285	83.285	83.644	83.644
Observations	918	918	933	933
	Female students in the task-based goals experiment		**Female students in the performance-based goals experiment**	
	Total points score		Total points score	
	OLS	Median	OLS	Median
Effect of asking students to set task-based goals	0.017	0.471		
	(0.637)	(0.652)		
	[1.000]	[0.941]		
Effect of asking students to set performance-based goals			-0.304	-0.810
			(0.625)	(0.689)
			[1.000]	[0.480]
Effect / (SD in Control group)	0.002	0.045	-0.029	-0.076
Mean of dependent variable in Control group	82.966	82.966	82.864	82.864
Observations	1,086	1,086	1,034	1,034

Notes: The regressions are the same as those reported in Table 5, except that we no longer include controls for student characteristics. Heteroskedasticity-consistent standard errors are shown in round brackets and two-sided Bonferonni-adjusted p-values are shown in square brackets. The Bonferonni adjustment accounts for the multiple null hypotheses being considered, i.e., zero treatment effect for men and zero treatment effect for women. *, ** and *** denote, respectively, significance at the 10%, 5% and 1% levels (two-sided tests based on the Bonferonni-adjusted p-values).

Table A.7: Gender differences in the effects of task-based goals and performance-based goals on student performance without controls for student characteristics

Students in the Control group of the task-based goals experiment			
	Points scored in one of the midterms or the final exam		
	(1)	(2)	(3)
	Male students	Female students	Male-Female difference
Completed practice exams	3.342	3.571	-0.229
	(0.475)	(0.443)	(0.650)
	[0.000]	[0.000]	[0.725]
Completed practice exams squared	-0.387	-0.386	-0.001
	(0.082)	(0.074)	(0.110)
	[0.000]	[0.000]	[0.995]
Average marginal effect	1.220	1.080	0.140
	(0.127)	(0.088)	(0.154)
	[0.000]	[0.000]	[0.363]
Mean of dependent variable in Control group	23.491	23.349	
Student fixed effects	Yes	Yes	
Observations (student-exam pairs)	1,389	1,671	

Notes: Columns (1) and (2) report fixed effects panel regressions of points scored in one of the midterms or the final exam on the number of practice exams completed in preparation for that midterm or final exam and the square of this variable. Each student-exam pair is an observation (giving three observations per student). We include a fixed effect for each student, and the fixed effects absorb any effects of student characteristics on student performance. The sample includes only students from the Control group of the task-based goals experiment, who were not asked to set goals. Average marginal effects are obtained as follows: for each student exam-pair in the Control group of the task-based goals experiment, we calculate the marginal effect of completing one extra practice exam on points scored in that midterm or final exam (the marginal effect for a student exam-pair is zero if the student completed the maximum number of practice exams for that exam); and we then average across all student exam-pairs in the Control group of the task-based goals experiment. Column (3) reports the male-female difference in the regression coefficients and the average marginal effect. Heteroskedasticity-consistent standard errors (with clustering at the student level) are shown in round brackets and two-sided p-values are shown in square brackets. Standard errors for the average marginal effect are calculated using the delta method. *, ** and *** denote, respectively, significance at the 10%, 5% and 1% levels (two-sided tests).

Table A.8: Estimates of the gender difference in marginal productivity

Web Appendix II Figures

Component	Points available	Points scored	Answer key
Syllabus Quiz	2	2	N/A
Start-Of-Course Survey	2	2	N/A
Quiz 1	3	2	Answer Key
Quiz 2	3	3	Answer Key
Quiz 3	3	2	Answer Key
Quiz 4	3	2	Answer Key
Quiz 5	3		
Quiz 6	3		
Quiz 7	3		
Quiz 8	3		
Quiz 9	3		
Quiz 10	3		
Best Midterm	30		
Midterm 1			
Midterm 2			
Final Exam	34		
End-Of-Course-Survey	2		
Total Points	**100**	**13**	

Grade Key

Total Points Scored (out of 100)	Letter Grade
91 and above	A
90 to 88	A-
87 to 86	B+
85 to 81	B
80 to 78	B-
77 to 76	C+
75 to 70	C
69 to 66	D
65 and below	E

Figure A.1: Example gradecard for a student in the Control group (Fall 2013 semester)

Syllabus quiz and start-of-course survey	
Syllabus quiz	2 points for completion
Consent form	For treated and control students
Start-of-course survey	Treated students set goal for letter grade in course 2 points for completion

Online quizzes
10 online quizzes throughout the semester Each scored from 0 to 3 points

Midterm exam 1
Scored from 0 to 30 points Only maximum of midterm 1 & 2 scores counts for letter grade

Midterm exam 2
Scored from 0 to 30 points Only maximum of midterm 1 & 2 scores counts for letter grade

Final exam
Scored from 0 to 34 points

End-of-course survey
2 points for completion

Figure A.2: Fall 2013 semester timeline

Consent Form for Cornell University Research Team Study on Course Performance

Before you start the survey, I want to tell you about a "Cornell University Research Team" that is conducting research to evaluate which factors contribute to good performance on this course.

Research Method
The team will use:
- Survey responses.
- Grades from this course.
- Information held by the [University name] registrar (e.g., admissions data, demographic information).

Confidentiality
- All the information will be made anonymous.
- This means that your name will never be seen by the Cornell University Research Team and will not be associated with the findings.

What you will be asked to do in this study
Nothing.

Risks
There are no risks to you.

Right to withdraw from the study
You have the right to withdraw from the study at any time during the semester. If you withdraw there will be no consequences for you; your academic standing, record, or relationship with the university will not be affected. Details of how to withdraw are available from the course webpage.

Who to contact if you have questions about the study:
Cornell Research Team: [curt@cornell.edu]
Full contact details are available from the course webpage.

Who to contact about your rights as a participant in this study:
Cornell Institutional Review Board, Ithaca NY. Email: irbhp@cornell.edu, phone: 607-255-5138; website: www.irb.cornell.edu. Concerns/complaints can also be anonymously reported through Ethicspoint (web: www.hotline.cornell.edu, phone (toll-free): 1-866-293-3077). Full contact details are available from the course webpage.

The Cornell University Research Team would be very grateful if you'd be willing to consent to your data being used in this study. Remember that your name will never be seen by the Research Team and there is nothing you need to do. (If you choose not to consent, you will still receive [1%][2%] towards your score for this course from completing the survey).

Yes, I consent | No, I don't consent

Figure A.3: Consent form

Syllabus quiz and start-of-course survey	
Syllabus quiz	1 point for completion
Consent form	For treated and control students
Start-of-course survey	1 point for completion

Online quizzes
9 online quizzes throughout the semester Each scored from 0 to 3 points

Mid-course survey 1
Treated students set goal for score in midterm exam 1 2 points for completion

Midterm exam 1
Scored from 0 to 30 points Only maximum of midterm 1 & 2 scores counts for letter grade

Mid-course survey 2
Treated students set goal for score in midterm exam 2 2 points for completion

Midterm exam 2
Scored from 0 to 30 points Only maximum of midterm 1 & 2 scores counts for letter grade

Mid-course survey 3
Treated students set goal for score in final exam 2 points for completion

Final exam
Scored from 0 to 34 points

End-of-course survey
1 point for completion

Figure A.4: Spring 2014 semester timeline

Please set a goal for your grade in this course.

Think carefully before setting your goal.

The professor and the TA will not see your goal. However, each time you get your quiz, midterm and final scores back, your gradecard will remind you of your goal.

My goal for this course is:

- ⊙ A
- ⊙ A−
- ⊙ B+
- ⊙ B
- ⊙ B−
- ⊙ C+
- ⊙ C
- ⊙ D
- ⊙ E
- ⊙ Prefer not to say

Figure A.5: Fall 2013 semester goal-setting question in start-of-course survey

Please set a goal for your score in the [Midterm 1][Midterm 2][Final] Exam.

Think carefully before setting your goal.

The professor and the TA will not see your goal. However, each time you get your quiz, midterm and final exam scores back, your gradecard will remind you of your goal.

My goal for my score in the [Midterm 1][Midterm 2][Final] Exam is:

[ ] out of [30][30][34]

⊙ Prefer not to say

Figure A.6: Spring 2014 semester goal-setting question in mid-course surveys

The goal that you set for this course is: X

(You set your goal in the start-of-course survey)

Component	Points available	Points scored	Answer key
Syllabus Quiz	2	2	N/A
Start-Of-Course Survey	2	2	N/A
Quiz 1	3	2	Answer Key
Quiz 2	3	3	Answer Key
Quiz 3	3	2	Answer Key
...	...	...	...

Figure A.7: Fall 2013 semester goal reminder on gradecard

Component	Points available	Points scored	Answer key
... ...	... ...	... ...	... ...
Best Midterm	30	24	
Midterm 1		22	**Your goal(*): X**
Midterm 2		24	**Your goal(*): Y**
Final Exam	34		**Your goal(*): Z**
End-Of-Course-Survey	1		
Total Points	**100**	**41**	

* You set this goal as part of a Mid-Course Survey

Grade Key

Total Points Scored (out of 100)	Letter Grade
91 and above	A
...	...

Figure A.8: Spring 2014 semester goal reminder on gradecard

{Treated and control students}

Practice Exams

There will be 5 practice exams for the [Midterm 1][Midterm 2][Final] Exam. Each practice exam will contain the same number of questions as the [Midterm 1][Midterm 2][Final] Exam.

The practice exams for the [Midterm 1][Midterm 2][Final] Exam will become active when Mid-Course Survey [1][2][3] closes. You will receive a reminder email at that time.

{Treated students only}

Question 6

Please set a goal for the number of practice exams that you will complete out of the 5 practice exams for the [Midterm 1][Midterm 2][Final] Exam.

Think carefully before setting your goal.

The professor and the TA will not see your goal. However, when you take the practice exams you will be reminded of your goal.

My goal is to complete

[ ] out of the 5 practice exams for the [Midterm 1][Midterm 2][Final] Exam

⊙ Prefer not to say

Figure A.9: Fall 2014 & Spring 2015 semesters practice exams information and goal-setting question in mid-course surveys

Dear Econ2023 Students,

The 5 practice exams for the [Midterm 1][Midterm 2][Final] Exam are now active.

{TREATED STUDENTS ONLY} Your goal is to complete X out of the 5 practice exams (you set this goal as part of Mid-Course Survey [1][2][3]).

To go to the practice exams, please go to the course webpage and follow the link.

Note that you will not receive any further reminders about the practice exams for the [Midterm 1][Midterm 2][Final] Exam and the practice exams will close when the [Midterm 1][Midterm 2][Final] Exam begins.

This is an automated email from the ECO 2023 system.

Figure A.10: Fall 2014 & Spring 2015 semesters practice exams reminder email

Practice Exams for the [Midterm 1][Midterm 2][Final] Exam

You have completed X out of the 5 practice exams for the [Midterm 1][Midterm 2][Final] Exam

{TREATED STUDENTS ONLY} **Your goal is to complete Z out of the 5 practice exams** (you set this goal as part of Mid-course Survey [1][2][3]).

Instructions:

- To take one of the practice exams, click on the link below.
- You can only take each practice exam once.
- There is no time limit.
- After answering each question, you will be given the correct answer.
- **This will be your only opportunity to see the correct answer.**
- You will not be able to go back to previous questions.

Practice Exam X+1

Figure A.11: Fall 2014 & Spring 2015 semesters practice exams introductory screen

Practice Exam X for the [Midterm 1][Midterm 2][Final] Exam

Your score was **Y** out of [30][30][34]

You have now completed X out of the 5 practice exams for the [Midterm 1][Midterm 2][Final] Exam

{TREATED STUDENTS ONLY} **Your goal is to complete Z out of the 5 practice exams** (you set this goal as part of Mid-course Survey [1][2][3]).

Return To Practice Exams Screen

Figure A.12: Fall 2014 & Spring 2015 semesters practice exams feedback screen

Web Appendix III Using theory to interpret our findings

Web Appendix III.1 Motivation

In this section we suggest some hypotheses for our findings. For each type of goal (performance- and task-based), our approach is to write down a simple model of goal setting and then use this to generate possible hypotheses. We acknowledge that these models are not the only ones that we could have used, but we are not aiming to test theory. Rather, we are using theory to guide the analysis and interpretation of our findings.

Our models build on Koch and Nafziger (2011) and are inspired by two key concepts in behavioral economics: present bias and loss aversion. The concept of present bias captures the idea that people lack control because they place a high weight on current utility (Strotz, 1956).[43] More specifically, a present-biased discounter places more weight on current utility relative to utility n periods in the future than she does on utility at future time t relative to utility at time $t+n$. This implies that present-biased discounters exhibit time inconsistency, since their time preferences at different dates are not consistent with one another. In the context of education, a present-biased student might set out to exert her preferred level of effort, but when the time comes to attend class or review for a test she might lack the self-control necessary to implement these plans.[44] Strotz (1956) and Pollak (1968) were the first to analyze how time-inconsistent agents make choices anticipating the different time preferences of their future selves. Building on this insight, Strotz (1956) noted that present-biased agents can mitigate their self-control problem by using commitment devices to bind their future self.[45]

The concept of loss aversion captures the idea that people dislike falling behind a salient reference point (Kahneman and Tversky, 1979).[46] In the context of education, a loss-averse student might work particularly hard in an attempt to achieve a salient reference point (e.g., a particular grade in her course).

Together, the literatures on present bias and loss aversion suggest that self-set goals might serve as an effective commitment device. Specifically, self-set goals might act as salient reference points, helping present-biased agents to mitigate their self-control problem and so steer their effort toward its optimal level. Indeed, Koch and Nafziger (2011) developed a model of goal setting based on this idea that we build on here. Unlike us, however, Koch and Nafziger (2011)

[43] Present bias has been proposed as an explanation for aspects of many behaviors such as addiction (Gruber and Kőszegi, 2001), early retirement (Diamond and Kőszegi, 2003), smoking (Khwaja et al., 2007), welfare program participation (Fang and Silverman, 2009) and credit card borrowing (Meier and Sprenger, 2010). See Dhami (2016) for a recent comprehensive survey of the literature on present bias.

[44] Under standard (i.e., exponential) discounting this self-control problem disappears.

[45] We provide examples of such commitment devices in Web Appendix V.2.

[46] Loss aversion has been proposed as a foundation of a number of phenomena such as the endowment effect (Kahneman et al., 1990), small-scale risk aversion (Rabin, 2000), the disposition effect (Genesove and Mayer, 2001), and the role of expectations in single-agent decision-making (Bell, 1985; Kőszegi and Rabin, 2006) and in strategic interactions (Gill and Stone, 2010; Gill and Prowse, 2012).

did not explore the effectiveness of different types of goals.[47]

Web Appendix III.2 Performance-based goal setting

Web Appendix III.2.1 Simple model

We start by building a simple model of performance-based goal setting. In Web Appendix III.2.2 we use the model to suggest three hypotheses for why performance-based goals might not be very effective in the context that we studied.

At period one the student sets a goal $g \geq 0$ for performance $f \geq 0$; we call the student at period one the *student-planner*. At period two the student chooses effort $e \geq 0$; we call the student at period two the *student-actor*. The student-actor incurs a cost of effort $C(e) = ce^2/2$, with $c > 0$. At period three performance is realized and the student incurs any disutility from failing to achieve her goal; we call the student at period three the *student-beneficiary*. Performance increases one-to-one in effort exerted by the student-actor at period two, i.e., $f(e) = e$, and the student-beneficiary's utility increases one-to-one in performance.[48] The student-beneficiary is loss averse around her goal: she suffers goal disutility that depends linearly on how far performance falls short of the goal set by the student-planner at period one. The student-beneficiary's goal disutility is given by $-l \max\{g - f(e), 0\}$. The parameter $l > 0$ captures the strength of loss aversion, which in our context we call the 'strength of goal disutility'.[49,50]

The student is present biased. In particular, the student exhibits quasi-hyperbolic discounting, with $\beta \in (0, 1)$ and $\delta \in (0, 1]$: the student discounts utility n periods in the future by a factor $\beta\delta^n$.[51] Under quasi-hyperbolic discounting the student-planner discounts period-two utility by a factor $\beta\delta$ and period-three utility by a factor $\beta\delta^2$, and so discounts period-three utility by δ relative to period-two utility. The student-actor, on the other hand, discounts period-three utility by $\beta\delta$ relative to immediate period-two utility. Since $\beta\delta < \delta$, the student-planner places more weight on utility from performance at period three relative to the cost of effort at period two than does the student-actor.

As a result of this present bias, and in the absence of a goal, the student-planner's desired effort is higher than the effort chosen by the student-actor: that is, the student exhibits a

[47] Koch and Nafziger (2011)'s model also differs from our models in that agents in their model choose from only two possible effort levels, while our models allow students to choose both effort and goals from a continuum. Again without exploring the effectiveness of different types of goals, Jain (2009) also studies theoretically how present-biased agents can use goals as reference points; in Jain (2009)'s model utility is discontinuous at the reference point, rather than kinked as in Kahneman and Tversky (1979)'s model of loss aversion that we and Koch and Nafziger (2011) use. Heath et al. (1999) and Wu et al. (2008) linked goals to loss aversion, but did not make the connection to present bias. Finally, a related theoretical literature studies expectations as goals (Suvorov and Van de Ven, 2008; Hsiaw, 2013; Hsiaw, 2016; Koch and Nafziger, 2016).

[48] Using one-to-one relationships instead of more general linear relationships is without loss of generality.

[49] Specifically, l measures the psychological loss from failing to achieve the goal relative to 'material' utility.

[50] This formulation implies that the student suffers disutility from failing to achieve her goal. However, it also implies that she enjoys no elation from exceeding the goal. This latter assumption can be justified in two ways. First, the more parsimonious one-parameter model of loss aversion allows us to gain useful insights into the effectiveness of goal setting. Second, if students enjoyed elation from exceeding their goals, they would have a strategic incentive to set low goals in order to enjoy the utility boost from exceeding them, but we do not see evidence that this motivation is an important driver of behavior in our data.

[51] Laibson (1997) was the first to apply the analytically tractable quasi-hyperbolic (or 'beta-delta') model of discounting to analyze the choices of present-biased time-inconsistent agents. Like us, Laibson (1997) finds the equilibria of a dynamic game among a sequence of temporal selves.

self-control problem due to time inconsistency.[52] We formalize this as follows:

Remark A.1

In the absence of a goal the student exhibits time inconsistency:
(i) The student-actor chooses effort $\underline{e} = \beta\delta/c$.
(ii) The student-planner would like the student-actor to exert effort $\hat{e} = \delta/c > \underline{e}$.

Proof. See Web Appendix IV.1. ■

To alleviate her self-control problem due to time-inconsistency, the student-planner might choose to set a goal. Goals can be effective by increasing the student-actor's marginal incentive to work in order to avoid the goal disutility that results from failing to achieve the goal.

To demonstrate this point, we solve for the subgame-perfect Nash equilibria of the game outlined above, in which the players are the student-planner and student-actor. We do so by backward induction. First, we analyze the effort choice of the student-actor at period two for any goal set by the student-planner at period one. The student-actor's utility is given by:

$$\begin{aligned} u_{act}(e|g) &= \beta\delta[f(e) - l\max\{g - f(e), 0\}] - C(e) && (1)\\ &= \beta\delta[e - l\max\{g - e, 0\}] - \frac{ce^2}{2}. && (2) \end{aligned}$$

Proposition A.1 shows how the student-actor's effort responds to the goal.

Proposition A.1

Let $\overline{e} = \beta\delta(1+l)/c$ *and recall from Remark A.1 that* $\underline{e} = \beta\delta/c < \overline{e}$ *denotes the student-actor's effort in the absence of a goal.*
(i) When $g \leq \underline{e}$, *the student-actor exerts effort* $e^* = \underline{e}$.
(ii) When $g \in [\underline{e}, \overline{e}]$, *the student-actor exerts effort* $e^* = g$.
(iii) When $g \geq \overline{e}$, *the student-actor exerts effort* $e^* = \overline{e}$.

Proof. See Web Appendix IV.1. ■

Proposition A.1 tells us that, perhaps unsurprisingly, the goal does not raise effort when it is set lower than the student-actor's optimal level of effort in the absence of a goal $\underline{e}$. Intermediate goals are effective: intermediate goals induce the student-actor to work hard enough to achieve the goal in order to avoid disutility from falling short of the goal. Beyond a certain point the marginal cost of effort outweighs the marginal reduction in goal disutility, and so the goal induces an increase in effort only to an upper bound $\overline{e}$. Goals above the upper bound leave the student-actor to suffer some goal disutility. This upper bound increases as the time-inconsistency problem becomes less severe (higher β) and as the strength of goal disutility l goes up.

Having established how the student-actor's effort responds to any goal set by the student-planner, we now consider the student-planner's optimal choice of goal. Letting $e^*(g)$ represent

[52] When $g = 0$, goal disutility is zero since $\max\{g - f(e), 0\} = 0$, and so $g = 0$ is equivalent to the absence of a goal.

the student-actor's optimal effort given a goal g, the student-planner's utility is given by:

$$u_{plan}(g|e^*(g)) = \beta\delta^2[f(e^*(g)) - l\max\{g - f(e^*(g)), 0\}] - \beta\delta C(e^*(g)) \quad (3)$$

$$= \beta\delta^2[e^*(g) - l\max\{g - e^*(g), 0\}] - \beta\delta\frac{c[e^*(g)]^2}{2}. \quad (4)$$

Proposition A.2

Recall from Remark A.1 that $\underline{e} = \beta\delta/c$ *and* $\hat{e} = \delta/c$ *denote, respectively, student-actor effort and student-planner desired effort in the absence of a goal.*

Recall from Proposition A.1 that $\overline{e} = \beta\delta(1 + l)/c$ *denotes maximal student-actor effort in the presence of a goal.*

(i) The optimal choice of goal for the student-planner is given by $g^* = \min\{\hat{e}, \overline{e}\}$*.*

(ii) When $\beta(1 + l) < 1$, $g^* = \overline{e}$*.*

(iii) When $\beta(1 + l) \geq 1$, $g^* = \hat{e}$*.*

(iv) Effort of the student-actor $e^* = g^* > \underline{e}$*, and so the student-actor works harder than in the absence of goal.*

Proof. See Web Appendix IV.1. ■

We know from Proposition A.1 that goals in the range $[\underline{e}, \overline{e}]$ induce the student-actor to work to achieve the goal, but that higher goals are ineffective in raising effort above $\overline{e}$. Thus, the student-planner will never set a goal higher than $\overline{e}$, since higher goals are not effective but leave the student-actor to suffer some goal disutility from failing to achieve the goal. If the student-planner could simply impose a level of effort on the student-actor, then we know from Remark 1 that she would choose $\hat{e}$. When $\beta(1+l)$ is big enough, $\hat{e} \leq \overline{e}$, and so the student-planner achieves her desired level of effort by setting $g^* = \hat{e}$. This case holds when the time-inconsistency problem is not too severe (high β) or the strength of goal disutility l is sufficiently high. When her desired level of effort is not achievable, the student-planner sets $g^* = \overline{e}$, and so the student-planner uses the goal to induce as much effort from the student-actor as she is able to. In either case, the optimal goal induces the student to work harder than she would in the absence of a goal and the student always achieves her goal in equilibrium.

Web Appendix III.2.2 Why might performance-based goals not be very effective?

This model of performance-based goal setting suggests that performance-based goals can improve course performance. However, our experimental data show that performance-based goals had a positive but small and statistically insignificant effect on student performance (Table 4). In our view, the simple model sketched above suggests three hypotheses for why performance-based goals might not be very effective in the context that we studied (we view these hypotheses as complementary).

Timing of goal disutility

In the simple model, the student works in period two and experiences any goal disutility from failing to achieve her performance-based goal in period three (i.e., when performance is realized).

This temporal distance will dampen the motivating effect of the goal. Formally, the student-actor discounts goal disutility by a factor $\beta\delta < 1$; in the expression for maximal student-actor effort in the presence of a goal, $\overline{e}$, the parameter measuring the strength of goal disutility, l, is multiplied by this discount factor (see Proposition A.2). Even when the temporal distance between effort and goal disutility is modest, the timing of goal disutility dampens the effectiveness of performance-based goals because quasi-hyperbolic discounters discount the near future relative to the present by a factor β even if $\delta \approx 1$ over the modest temporal distance.

Overconfidence

In the simple model, students understand perfectly the relationship between effort and performance. In contrast, the education literature suggests that students face considerable uncertainty about the educational production function, and that this uncertainty could lead to students holding incorrect beliefs about the relationship between effort and performance (e.g., Romer, 1993, and Fryer, 2013). Furthermore, the broader behavioral literature shows that people tend to be overconfident when they face uncertainty (e.g., Weinstein, 1980, Camerer and Lovallo, 1999, and Park and Santos-Pinto, 2010). The behavioral literature also provides a number of theoretical underpinnings for overconfidence (e.g., Brunnermeier and Parker, 2005, Johnson and Fowler, 2011, and Gossner and Steiner, 2016).

Suppose that some students are overconfident. An overconfident student believes that the production function is given by $f(\cdot)$, when in fact performance is given by $hf(\cdot)$ with $h \in (0, 1)$ for any value of effort. An overconfident student will act as if the production function is given by $f(\cdot)$, and so the model in Web Appendix III.2.1 describes her choice of goal and effort. However, the student's actual performance with goal setting and in the absence of a goal will be a proportion h of that expected by the student. As a result, the impact of performance-based goal setting on performance will be reduced by this proportion h. Furthermore, the overconfident student will unexpectedly fail to achieve her performance-based goal.

Performance uncertainty

In the simple model, the student knows for sure how her effort translates into performance (i.e., the relationship between effort and performance involves no uncertainty). As such, her goal is always achieved in equilibrium. In practice, the relationship between effort and performance is likely to be noisy and, as in our experiment, performance-based goals are not always reached. The student could face uncertainty about her own ability or about the productivity of work effort. The student might also get unlucky: for instance, the draw of questions on the exam might be unfavorable or the student might get sick near the exam.

To introduce uncertainty about performance in a straightforward way, suppose that the student is risk neutral and that with known probability $\pi \in (0, 1)$ her effort translates into performance according to $f(\cdot)$ as in Web Appendix III.2.1, while with probability $1 - \pi$ performance $f = 0$ (since we assume that π is known, the student is neither overconfident nor underconfident).[53]

[53]We can think of $f = 0$ as a baseline level of performance that the student achieves with little effort even in the absence of goal setting.

The formal details of the analysis are relegated to Web Appendix IV.3. The uncertainty directly reduces the student-actor's marginal incentive to exert effort, which reduces by a factor π the equilibrium goal and equilibrium effort with and without goal setting. However, this general reduction in incentives is not the only effect of uncertainty: performance-based goals become risky because when performance turns out to be low the student fails to achieve her performance-based goal and so suffers goal disutility that increases in the goal (as in the simple model, goals are never exceeded in equilibrium). In contrast to the case of overconfidence discussed above, goal failure is not unexpected: the student facing uncertainty anticipates that she will not always achieve her performance-based goal.

Anticipating the goal disutility suffered when performance turns out to be low, the student-planner further scales back the performance-based goal that she sets for the student-actor, which reduces the effectiveness of performance-based goal setting. Formally, this extra effect adds the second term to the numerator in the expression for $\tilde{e}$ in Proposition A.6 in Web Appendix IV.3.[54,55]

Web Appendix III.3 Task-based goal setting

Web Appendix III.3.1 Simple model

We now model task-based goal setting. At period one the student-planner sets a goal $g \geq 0$ for the number of units of the task to complete $a \geq 0$. We call a the 'level of task completion'. At period two the student-actor chooses the level of task completion a. The student-actor incurs a cost of task completion $C(a) = \kappa a^2/2$, with $\kappa > 0$. Furthermore, at period two the loss-averse student-actor suffers goal disutility that depends linearly on how far the level of task completion falls short of the goal set by the student-planner at period one: she suffers goal disutility of $-\lambda \max\{g - a, 0\}$, where the loss parameter λ captures the strength of goal disutility.[56] At period three performance is realized. Performance increases linearly in the level of task completion: $f(a) = \theta a$, with $\theta > 0$; while the student-beneficiary's utility increases one-to-one in performance.[57] The present-biased student exhibits quasi-hyperbolic discounting

[54] Proposition A.6 in Web Appendix IV.3 focuses on the case in which uncertainty is not too big. When the student faces a lot of uncertainty, the extra effect could lead the student-planner to prefer not to set a goal at all.

[55] This scaling back of goals is not necessarily at odds with the fact that the performance-based goals that we see in the data appear ambitious. First, the goal will appear ambitious relative to average achievement because, as noted above, when performance turns out to be low the student fails to achieve her goal. Second, without any scaling back the goals might have been even higher. Third, the overconfidence that we discuss above could keep the scaled-back goal high. Fourth, we explain in Web Appendix III.3.4 below that students likely report as their goal an 'aspiration' that is only relevant if, when the time comes to study, the cost of effort turns out to be particularly low: the actual cost-specific goal that the student aims to hit could be much lower than this aspiration.

[56] As explained in footnote 50, we do not include any elation from exceeding the goal. We use new notation for the parameter that measures the strength of goal disutility λ because the units used to measure the level of task completion are different from the performance units in Web Appendix III.2.1. Note also that goal disutility is incurred at period two here because the student-actor observes how far she is from the task-based goal immediately when she stops working on the task. For performance-based goals in Web Appendix III.2.1 goal disutility is incurred at period three when performance is realized.

[57] Units of performance and units of task completion are both defined externally, and so we need to introduce the parameter θ to model a linear relationship between them.

as described in Web Appendix III.2.1. Thus the student-actor's utility is given by:

$$u_{act}(a|g) = \beta\delta f(a) - [\lambda \max\{g-a, 0\} + C(a)] \quad (5)$$

$$= \beta\delta\theta a - \left[\lambda \max\{g-a, 0\} + \frac{\kappa a^2}{2}\right]; \quad (6)$$

and, letting $a^*(g)$ represent the student-actor's optimal level of task completion given a goal g, the student-planner's utility is given by:

$$u_{plan}(g|a^*(g)) = \beta\delta^2 f(a^*(g)) - \beta\delta[\lambda \max\{g - a^*(g), 0\} + C(a^*(g))] \quad (7)$$

$$= \beta\delta^2\theta a^*(g) - \beta\delta\left[\lambda \max\{g - a^*(g), 0\} + \frac{\kappa[a^*(g)]^2}{2}\right]. \quad (8)$$

When we solve the game by backward induction, we get results that are qualitatively similar to those for performance-based goals in Web Appendix III.2.1. The formal results and proofs are relegated to Web Appendix IV.2. The three relevant thresholds now become:

$$\underline{a} = \frac{\beta\delta\theta}{\kappa}; \qquad \hat{a} = \frac{\delta\theta}{\kappa} > \underline{a}; \qquad \overline{a} = \frac{\beta\delta\theta + \lambda}{\kappa} > \underline{a}. \quad (9)$$

Mirroring Remark A.1, in the absence of a goal the present-biased student exhibits a self-control problem due to time inconsistency: the student-actor chooses a level of task completion $\underline{a}$, which is smaller than the student-planner's desired level of task completion $\hat{a}$. The upper bound on student-actor task completion in the presence of a goal is given by $\overline{a}$. Mirroring Proposition A.1, this upper bound increases as the time-inconsistency problem becomes less severe (higher β) and as the strength of goal disutility λ goes up. Mirroring Proposition A.2, the optimal choice of goal for the student-planner is given by $g^* = \min\{\hat{a}, \overline{a}\}$ and the optimal goal induces a level of task completion by the student-actor of $a^* = g^* > \underline{a}$; the optimal goal induces a higher level of task completion than in the absence of a goal, and the student always achieves her goal in equilibrium.[58] The goal increases the level of task completion as well as improving course performance.

Web Appendix III.3.2 Why were task-based goals effective?

Our experimental data show that task-based goals improved task completion and course performance (see Table 2 for the effect on task completion and Table 4 for the effect on course performance).[59] How might we account for these findings, given our analysis of why performance-based goals might not be very effective? In our view, an obvious answer is that with task-based goal setting, the three factors that reduce the effectiveness of performance-based goals (Web Appendix III.2.2) are of lesser importance or do not apply at all.

[58] When $\beta\delta\theta + \lambda \geq \delta\theta$, $\hat{a} \leq \overline{a}$, and so the student-planner achieves her desired level of task completion by setting $g^* = \hat{a}$. Similarly to Proposition A.2, this case holds when the time-inconsistency problem is not too severe (high β) or the strength of goal disutility λ is sufficiently high.

[59] It is possible that some students in the Control group (who were not invited to set goals) might already use goals as a commitment device. However, since we find that task-based goals are successful at increasing performance, we conclude that many students in the Control group did not use goals or set goals that were not fully effective. We note that asking students to set goals might make the usefulness of goal setting as a commitment device more salient and so effective. Reminding students of their goal, as we did, might also help to make them more effective.

Timing of goal disutility

In the case of task-based goal setting, any goal disutility from failing to achieve the task-based goal is suffered immediately when the student stops working on the task in period two. Thus, unlike the case of performance-based goal setting discussed in Web Appendix III.2.2, there is no temporal distance that dampens the motivating effect of the goal. Formally, in the expression for maximal task completion in the presence of a goal, $\overline{a}$, the parameter measuring the strength of goal disutility, λ, is undiscounted (see (9)).

Overconfidence

As discussed in Web Appendix III.2.2, overconfidence reduces the effectiveness of goal setting. Recall that an overconfident student acts as if the production function is given $f(\cdot)$, when in fact performance is given by $hf(\cdot)$, which reduces the impact of goal setting on performance by the proportion h. In the case of task-based goal setting, this effect is mitigated if practice exams direct students toward productive tasks: in that case h goes up. Plausibly, teachers have better information about which tasks are likely to be productive, and asking students to set goals for productive tasks is one way to improve the power of goal setting for overconfident students.

Instead of improving the power of goal setting by directing overconfident students toward productive tasks, it is conceivable that task-based goals improved performance via another channel: signaling to students in the Treatment group that practice exams were an effective task. But we think this is highly unlikely. First, we were careful to make the practice exams as salient as possible to the Control group. Second, students in the Control group in fact completed many practice exams. Third, it is hard to understand why only men would respond to the signal.

Performance uncertainty

In Web Appendix III.2.2, we introduced uncertainty about performance. It is straightforward to add performance uncertainty into the simple model of task-based goal setting outlined in Web Appendix III.3.1. The formal details of the analysis are relegated to Web Appendix IV.3. The important point to note is that even with uncertainty about performance, the student continues to achieve her task-based goal: there is no 'task uncertainty'. The reason is that the student-actor controls the number of units of the task that she completes and so can guarantee to hit her task-based goal. Thus, unlike the case of performance-based goals with uncertainty, the student has no reason to scale back her task-based goal to reduce goal disutility in the event that the goal is not reached.

Web Appendix III.3.3 Why were task-based goals more effective for men than for women?

Our data show that task-based goals are more effective for men than for women. More specifically: in the Control group without goal setting men completed fewer practice exams than women (Table 3); and task-based goals increased performance and the number of practice exams completed more for men than for women (Tables 5 and 3 and respectively). In the context of our simple model of task-based goal setting (Web Appendix III.3.1), one hypothesis that could explain this finding is that the male students in our sample are more present biased than the female students (i.e., the men have a lower β parameter). Existing empirical evidence supports the idea that men may have less self-control and be more present biased than women (see Web Appendix V.6 for a survey of this evidence.) Interestingly, in a laboratory experiment in which goals were set by the experimenter rather than by the subjects themselves, Smithers (2015) finds that goals increased the work performance of men but not that of women.

To understand the role of present bias, first note that the student-actor's level of task completion in the absence of a goal $\underline{a}$ is increasing in β: the more present biased the student, the fewer practice exams he or she completes without a goal. Thus, if men are more present biased than women, then their higher degree of present bias will push down their level of task completion in the Control group relative to that of women. Second, the increase in task completion induced by goal setting also depends on the degree of present bias: in particular, the difference between the student-planner's desired level of task completion $\hat{a}$ and the student-actor's level of task completion in the absence of a goal $\underline{a}$ is decreasing in β. Thus, if men are more present biased than women, then goal setting will tend to be more effective at increasing the number of practice exams that men complete, which in turn feeds into a larger effect on performance.

Web Appendix III.3.4 Why were task-based goals not always achieved?

In the simple models outlined in Web Appendix III.2.1 and Web Appendix III.3.1 goals are always achieved. In Web Appendix III.2.2 we explained how overconfidence and performance uncertainty can result in a student's failure to achieve her performance-based goal. A puzzle remains: even though task-based goals are more frequently achieved than performance-based goals, task-based goals are not always achieved (Table 1).

Failure to achieve task-based goals emerges naturally if we relax the assumption that costs are known with certainty. In particular, suppose that the student-actor's cost parameter (c or κ) can be high or low, and that the student-actor draws her cost parameter in period two before she decides how hard to work (the analysis extends naturally to any number of possible cost draws). For example, the cost uncertainty could be driven by uncertainty about the set of leisure activities available to the student during the time that she planned to study. Anticipating this cost uncertainty, we allow the student-planner to set a goal for both possible cost draws. The optimal goal for a given cost draw is just as in the models in Web Appendix III.2.1 and Web Appendix III.3.1 with no cost uncertainty, and the student-actor always works hard enough to achieve the cost-specific goal. Of course, we ask the student to report only one goal: we assume here that the student-planner reports to us only her goal for the low-cost draw. This goal is like an aspiration: if the cost turns out to be high, the goal is scaled down to reflect the higher cost. Because we as the experimenter observe only the reported aspiration, when the cost is high we observe a failure to achieve the reported aspiration, even though the student achieves her cost-specific goal.

Web Appendix IV Formal results and proofs

Web Appendix IV.1 Proofs for Web Appendix III.2.1

Proof of Remark A.1.

In the absence of a goal, the student-actor's utility and the student-planner's utility are given by, respectively:

$$\begin{aligned} u_{act}(e) &= \beta\delta f(e) - C(e) && (10)\\ &= \beta\delta e - \frac{ce^2}{2}; && (11)\\ u_{plan}(e) &= \beta\delta^2 f(e) - \beta\delta C(e) && (12)\\ &= \beta\delta^2 e - \beta\delta\frac{ce^2}{2}. && (13) \end{aligned}$$

Both utilities are strictly concave since $c > 0$. Straightforward maximization then gives the result. ■

Proof of Proposition A.1.

Using (2), on the range $e \in [0, g]$:

$$\begin{aligned} \frac{\partial u_{act}}{\partial e} &= \beta\delta(1 + l) - ce; && (14)\\ \frac{\partial^2 u_{act}}{\partial e^2} &= -c < 0; \quad \text{and so} && (15)\\ e^* &= \min\{\overline{e}, g\}. && (16) \end{aligned}$$

Using (2), on the range $e \in [g, \infty)$:

$$\begin{aligned} \frac{\partial u_{act}}{\partial e} &= \beta\delta - ce; && (17)\\ \frac{\partial^2 u_{act}}{\partial e^2} &= -c < 0; \quad \text{and so} && (18)\\ e^* &= \max\{\underline{e}, g\}. && (19) \end{aligned}$$

(i) When $g \leq \underline{e}$, on the range $e \in [0, g]$, $e^* = g$, and on the range $e \in [g, \infty)$, $e^* = \underline{e}$. Thus, on the range $e \in [0, \infty)$, $e^* = \underline{e}$.

(ii) When $g \in [\underline{e}, \overline{e}]$, on the range $e \in [0, g]$, $e^* = g$, and on the range $e \in [g, \infty)$, $e^* = g$. Thus, on the range $e \in [0, \infty)$, $e^* = g$.

(iii) When $g \geq \overline{e}$, on the range $e \in [0, g]$, $e^* = \overline{e}$, and on the range $e \in [g, \infty)$, $e^* = g$. Thus, on the range $e \in [0, \infty)$, $e^* = \overline{e}$. ■

Proof of Proposition A.2.

On the range $g \in [0, \underline{e}]$, $e^*(g) = \underline{e}$ from Proposition A.1, and so $\partial e^*/\partial g = 0$ and $\max\{g - e^*(g), 0\} = 0$. Using (4), $du_{plan}/dg = 0$ and so any $g \in [0, \underline{e}]$ is optimal (including $\underline{e}$).

On the range $g \in [\underline{e}, \overline{e}]$, $e^*(g) = g$ from Proposition A.1, and so $\partial e^*/\partial g = 1$ and

$\max\{g - e^*(g), 0\} = 0$. Using (4), and noting that $\hat{e} > \underline{e}$ and $\overline{e} > \underline{e}$:

$$\frac{du_{plan}}{dg} = \beta\delta^2 - \beta\delta c g; \quad (20)$$

$$\frac{d^2u_{plan}}{dg^2} = -\beta\delta c < 0; \quad \text{and so} \quad (21)$$

$$g^* = \min\{\hat{e}, \overline{e}\} > \underline{e}. \quad (22)$$

On the range $g \in [\overline{e}, \infty)$, $e^*(g) = \overline{e}$ from Proposition A.1, and so $\partial e^*/\partial g = 0$ and $\max\{g - e^*(g), 0\} = g - \overline{e}$. Using (4), $du_{plan}/dg = -\beta\delta^2 l < 0$ and so $g^* = \overline{e}$.

Stitching the ranges together gives $g^* = \min\{\hat{e}, \overline{e}\} > \underline{e}$. Parts (ii) and (iii) follow, given that $\overline{e} < \hat{e} \Leftrightarrow \beta(1 + l) < 1$. Finally, (iv) follows immediately from Proposition A.1. ■

Web Appendix IV.2 Results and proofs for Web Appendix III.3.1

Remark A.2

In the absence of a goal the student exhibits time inconsistency:
(i) The student-actor chooses a level of task completion $\underline{a} = \beta\delta\theta/\kappa$.
(ii) The student-planner would like the student-actor to choose a level of task completion $\hat{a} = \delta\theta/\kappa > \underline{a}$.

Proof of Remark A.2.

In the absence of a goal, the student-actor's utility and the student-planner's utility are given by, respectively:

$$\begin{aligned} u_{act}(a) &= \beta\delta f(a) - C(a) && (23)\\ &= \beta\delta\theta a - \frac{\kappa a^2}{2}; && (24)\\ u_{plan}(a) &= \beta\delta^2 f(a) - \beta\delta C(a) && (25)\\ &= \beta\delta^2\theta a - \beta\delta\frac{\kappa a^2}{2}. && (26) \end{aligned}$$

Both utilities are strictly concave since $\kappa > 0$. Straightforward maximization then gives the result. ■

Proposition A.3

Let $\overline{a} = (\beta\delta\theta + \lambda)/\kappa$ *and recall from Remark A.2 that* $\underline{a} = \beta\delta\theta/\kappa < \overline{a}$ *denotes the student-actor's level of task completion in the absence of a goal.*
(i) When $g \leq \underline{a}$, *the student-actor chooses a level of task completion* $a^* = \underline{a}$.
(ii) When $g \in [\underline{a}, \overline{a}]$, *the student-actor chooses a level of task completion* $a^* = g$.
(iii) When $g \geq \overline{a}$, *the student-actor chooses a level of task completion* $a^* = \overline{a}$.

Proof of Proposition A.3.

Using (6), on the range $a \in [0, g]$:

$$\begin{aligned} \frac{\partial u_{act}}{\partial a} &= \beta\delta\theta + \lambda - \kappa a; && (27)\\ \frac{\partial^2 u_{act}}{\partial a^2} &= -\kappa < 0; \quad \text{and so} && (28)\\ a^* &= \min\{\overline{a}, g\}. && (29) \end{aligned}$$

Using (6), on the range $a \in [g, \infty)$:

$$\begin{aligned} \frac{\partial u_{act}}{\partial a} &= \beta\delta\theta - \kappa a; && (30)\\ \frac{\partial^2 u_{act}}{\partial a^2} &= -\kappa < 0; \quad \text{and so} && (31)\\ a^* &= \max\{\underline{a}, g\}. && (32) \end{aligned}$$

(i) When $g \leq \underline{a}$, on the range $a \in [0, g]$, $a^* = g$, and on the range $a \in [g, \infty)$, $a^* = \underline{a}$. Thus, on the range $a \in [0, \infty)$, $a^* = \underline{a}$.

(ii) When $g \in [\underline{a}, \overline{a}]$, on the range $a \in [0, g]$, $a^* = g$, and on the range $a \in [g, \infty)$, $a^* = g$. Thus, on the range $a \in [0, \infty)$, $a^* = g$.

(iii) When $g \geq \overline{a}$, on the range $a \in [0, g]$, $a^* = \overline{a}$, and on the range $a \in [g, \infty)$, $a^* = g$. Thus, on the range $a \in [0, \infty)$, $a^* = \overline{a}$. ■

Proposition A.4

Recall from Remark A.2 that $\underline{a} = \beta\delta\theta/\kappa$ and $\hat{a} = \delta\theta/\kappa$ denote, respectively, student-actor task completion and student-planner desired task completion in the absence of a goal.
Recall from Proposition A.3 that $\overline{a} = (\beta\delta\theta + \lambda)/\kappa$ denotes maximal student-actor task completion in the presence of a goal.
(i) The optimal choice of goal for the student-planner is given by $g^* = \min\{\hat{a}, \overline{a}\}$.
(ii) When $\beta\delta\theta + \lambda < \delta\theta$, $g^* = \overline{a}$.
(iii) When $\beta\delta\theta + \lambda \geq \delta\theta$, $g^* = \hat{a}$.
(iv) Student-actor task completion $a^* = g^* > \underline{a}$, and so the level of task completion is higher than in the absence of goal.

Proof of Proposition A.4.

On the range $g \in [0, \underline{a}]$, $a^*(g) = \underline{a}$ from Proposition A.3, and so $\partial a^*/\partial g = 0$ and $\max\{g - a^*(g), 0\} = 0$. Using (8), $du_{plan}/dg = 0$ and so any $g \in [0, \underline{a}]$ is optimal (including $\underline{a}$).

On the range $g \in [\underline{a}, \overline{a}]$, $a^*(g) = g$ from Proposition A.3, and so $\partial a^*/\partial g = 1$ and $\max\{g - a^*(g), 0\} = 0$. Using (8), and noting that $\hat{a} > \underline{a}$ and $\overline{a} > \underline{a}$:

$$\frac{du_{plan}}{dg} = \beta\delta^2\theta - \beta\delta\kappa g; \tag{33}$$

$$\frac{d^2u_{plan}}{dg^2} = -\beta\delta\kappa < 0; \quad \text{and so} \tag{34}$$

$$g^* = \min\{\hat{a}, \overline{a}\} > \underline{a}. \tag{35}$$

On the range $g \in [\overline{a}, \infty)$, $a^*(g) = \overline{a}$ from Proposition A.3, and so $\partial a^*/\partial g = 0$ and $\max\{g - a^*(g), 0\} = g - \overline{a}$. Using (8), $du_{plan}/dg = -\beta\delta\lambda < 0$ and so $g^* = \overline{a}$.

Stitching the ranges together gives $g^* = \min\{\hat{a}, \overline{a}\} > \underline{a}$. Parts (ii) and (iii) follow, given that $\overline{a} < \hat{a} \Leftrightarrow \beta\delta\theta + \lambda < \delta\theta$. Finally, (iv) follows immediately from Proposition A.3 ■

Web Appendix IV.3 Performance uncertainty

Web Appendix IV.3, Part A: Task-based goals (see Web Appendix III.3.2)

For task-based goals, when we add uncertainty as described in the third part of Web Appendix III.3.2 the student-actor's utility (given by (5) and (6) with no uncertainty) and the student-planner's utility (given by (7) and (8) with no uncertainty) become, respectively:

$$Eu_{act}(a|g) = \beta\delta[\pi f(a) + (1-\pi)(0)] - [\lambda \max\{g-a, 0\} + C(a)] \tag{36}$$

$$= \beta\delta\pi\theta a - \left[\lambda \max\{g-a, 0\} + \frac{\kappa a^2}{2}\right]; \tag{37}$$

$$Eu_{plan}(g|a^*(g)) = \beta\delta^2[\pi f(a^*(g)) + (1-\pi)(0)] - \beta\delta[\lambda \max\{g - a^*(g), 0\} + C(a^*(g))] \tag{38}$$

$$= \beta\delta^2\pi\theta a^*(g) - \beta\delta \left[\lambda \max\{g - a^*(g), 0\} + \frac{\kappa[a^*(g)]^2}{2}\right]. \tag{39}$$

Comparing (37) and (39) to (6) and (8) without uncertainty, the only difference is that every θ in the case without uncertainty has been replaced by $\pi\theta$. Since $\theta \in (0,1)$ and $\pi\theta \in (0,1)$, the results for task-based goals without uncertainty described in Web Appendix IV.2 continue to hold with uncertainty, replacing every θ with $\pi\theta$.

Web Appendix IV.3, Part B: Performance-based goals (see Web Appendix III.2.2)

For performance-based goals, when we add uncertainty as described in the third part of Web Appendix III.2.2 the student-actor's utility (given by (1) and (2) with no uncertainty) and the student-planner's utility (given by (3) and (4) with no uncertainty) become, respectively:

$$Eu_{act}(e|g) = \beta\delta\{\pi[f(e) - l\max\{g - f(e), 0\}] + (1-\pi)[0 - l\max\{g-0, 0\}]\} - C(e) \tag{40}$$

$$= \beta\delta\{\pi[e - l\max\{g-e, 0\}] - (1-\pi)lg\} - \frac{ce^2}{2}; \tag{41}$$

$$Eu_{plan}(g|e^*(g)) =$$

$$\beta\delta^2\{\pi[f(e^*(g)) - l\max\{g - f(e^*(g)), 0\}] + (1-\pi)[0 - l\max\{g-0, 0\}]\} - \beta\delta C(e^*(g)) \tag{42}$$

$$= \beta\delta^2\{\pi[e^*(g) - l\max\{g - e^*(g), 0\}] - (1-\pi)lg\} - \beta\delta\frac{c[e^*(g)]^2}{2}. \tag{43}$$

Remark A.3

In the absence of a goal the student exhibits time inconsistency:

(i) The student-actor chooses effort $\underline{e} = \beta\delta\pi/c$.

(ii) The student-planner would like the student-actor to exert effort $\hat{e} = \delta\pi/c > \underline{e}$.

Proof of Remark A.3.

In the absence of a goal, the student-actor's utility and the student-planner's utility are given by, respectively:

$$\begin{aligned} Eu_{act}(e) &= \beta\delta[\pi f(e) + (1-\pi)(0)] - C(e) && (44) \\ &= \beta\delta\pi e - \frac{ce^2}{2}; && (45) \\ Eu_{plan}(e) &= \beta\delta^2[\pi f(e) + (1-\pi)(0)] - \beta\delta C(e) && (46) \\ &= \beta\delta^2\pi e - \beta\delta\frac{ce^2}{2}. && (47) \end{aligned}$$

Both utilities are strictly concave since $c > 0$. Straigthforward maximization then gives the result. ■

Proposition A.5

Let $\overline{e} = \beta\delta\pi(1+l)/c$ and recall from Remark A.3 that $\underline{e} = \beta\delta\pi/c < \overline{e}$ denotes the student-actor's effort in the absence of a goal.

(i) When $g \leq \underline{e}$, the student-actor exerts effort $e^* = \underline{e}$.

(ii) When $g \in [\underline{e}, \overline{e}]$, the student-actor exerts effort $e^* = g$.

(iii) When $g \geq \overline{e}$, the student-actor exerts effort $e^* = \overline{e}$.

Proof of Proposition A.5.

Using (41), on the range $e \in [0, g]$:

$$\begin{aligned} \frac{\partial Eu_{act}}{\partial e} &= \beta\delta\pi(1+l) - ce; && (48) \\ \frac{\partial^2 Eu_{act}}{\partial e^2} &= -c < 0; \quad \text{and so} && (49) \\ e^* &= \min\{\overline{e}, g\}. && (50) \end{aligned}$$

Using (41), on the range $e \in [g, \infty)$:

$$\begin{aligned} \frac{\partial Eu_{act}}{\partial e} &= \beta\delta\pi - ce; && (51) \\ \frac{\partial^2 Eu_{act}}{\partial e^2} &= -c < 0; \quad \text{and so} && (52) \\ e^* &= \max\{\underline{e}, g\}. && (53) \end{aligned}$$

(i) When $g \leq \underline{e}$, on the range $e \in [0, g]$, $e^* = g$, and on the range $e \in [g, \infty)$, $e^* = \underline{e}$. Thus, on the range $e \in [0, \infty)$, $e^* = \underline{e}$.

(ii) When $g \in [\underline{e}, \overline{e}]$, on the range $e \in [0, g]$, $e^* = g$, and on the range $e \in [g, \infty)$, $e^* = g$. Thus, on the range $e \in [0, \infty)$, $e^* = g$.

(iii) When $g \geq \overline{e}$, on the range $e \in [0, g]$, $e^* = \overline{e}$, and on the range $e \in [g, \infty)$, $e^* = g$. Thus, on the range $e \in [0, \infty)$, $e^* = \overline{e}$. ■

Proposition A.6

Recall from Remark A.3 that $\underline{e} = \beta\delta\pi/c$ *denotes student-actor effort in the absence of a goal. Recall from Proposition A.5 that* $\overline{e} = \beta\delta\pi(1+l)/c$ *denotes maximal student-actor effort in the presence of a goal.*

Let $\tilde{e} = [\delta\pi - \delta(1-\pi)l]/c$ *and recall from Remark A.3 that* $\hat{e} = \delta\pi/c > \tilde{e}$ *denotes the student-planner's desired effort in the absence of a goal.*

There exists a $\overline{\pi} \in (0,1)$ *such that for all* $\pi \in [\overline{\pi}, 1)$*:*

(i) The optimal choice of goal for the student-planner is given by $g^* = \min\{\tilde{e}, \overline{e}\}$*.*

(ii) When $\beta(1+l) + l(1-\pi)/\pi < 1$*,* $g^* = \overline{e}$*.*

(iii) When $\beta(1+l) + l(1-\pi)/\pi \geq 1$*,* $g^* = \tilde{e}$*.*

(iv) Effort of the student-actor $e^* = g^* > \underline{e}$*, and so the student-actor works harder than in the absence of goal.*

Proof of Proposition A.6.

On the range $g \in [0, \underline{e}]$, $e^*(g) = \underline{e}$ from Proposition A.5, and so $\partial e^*/\partial g = 0$ and $\max\{g - e^*(g), 0\} = 0$. Using (43), $dEu_{plan}/dg = -\beta\delta^2(1-\pi)l < 0$ and so $g^* = 0$. Note also that $\lim_{\pi\to 1}(dEu_{plan}/dg) = 0$, and so:

$$\lim_{\pi\to 1}\left[Eu_{plan}(g = \underline{e}) - Eu_{plan}(g = 0)\right] = 0. \tag{54}$$

On the range $g \in [\underline{e}, \overline{e}]$, $e^*(g) = g$ from Proposition A.5, and so $\partial e^*/\partial g = 1$ and $\max\{g - e^*(g), 0\} = 0$. Using (43), and noting that $\overline{e} > \underline{e}$:

$$\frac{dEu_{plan}}{dg} = \beta\delta^2[\pi - (1-\pi)l] - \beta\delta c g; \tag{55}$$

$$\frac{d^2Eu_{plan}}{dg^2} = -\beta\delta c < 0; \quad \text{and so} \tag{56}$$

$$g^* = \min\{\max\{\tilde{e}, \underline{e}\}, \overline{e}\}. \tag{57}$$

Note also that $\lim_{\pi\to 1}(\tilde{e} - \underline{e}) > 0$ and $\lim_{\pi\to 1}(\overline{e} - \underline{e}) > 0$. Thus, on the range $g \in [\underline{e}, \overline{e}]$, $g^* = \min\{\tilde{e}, \overline{e}\} > \underline{e}$ for π sufficiently close to 1.

When $\pi = 1$, from the proof of Proposition A.2, $g^* = \min\{\hat{e}, \overline{e}\}$ gives the student-planner strictly more utility than $g = \underline{e}$. Furthermore, $\lim_{\pi\to 1}(55) = (20)$, $\lim_{\pi\to 1}\tilde{e} = \hat{e}_{|\pi=1}$, $\lim_{\pi\to 1}\overline{e} = \overline{e}_{|\pi=1}$ and $\lim_{\pi\to 1}\underline{e} = \underline{e}_{|\pi=1}$. Thus:

$$\lim_{\pi\to 1}\left[Eu_{plan}(g = \min\{\tilde{e}, \overline{e}\}) - Eu_{plan}(g = \underline{e})\right] > 0. \tag{58}$$

On the range $g \in [\overline{e}, \infty)$, $e^*(g) = \overline{e}$ from Proposition A.5, and so $\partial e^*/\partial g = 0$ and $\max\{g - e^*(g), 0\} = g - \overline{e}$. Using (43), $dEu_{plan}/dg = -\beta\delta^2 l < 0$ and so $g^* = \overline{e}$.

Stitching the ranges together, and using (54) and (58), there exists a $\overline{\pi} \in (0,1)$ such that $g^* = \min\{\tilde{e}, \overline{e}\} > \underline{e}$ for all $\pi \in [\overline{\pi}, 1)$. Parts (ii) and (iii) follow, given that $\overline{e} < \tilde{e} \Leftrightarrow \beta(1+l) + l(1-\pi)/\pi < 1$. Finally, (iv) follows immediately from Proposition A.5. ■

Web Appendix V More details regarding related literature

Web Appendix V.1 Details referred to in footnote 1

Studies using randomized experiments and natural experiments to evaluate the effects of financial incentives on the performance of college students have been inconclusive: Henry et al. (2004), Cha and Patel (2010), Scott-Clayton (2011), De Paola et al. (2012) and Castleman (2014) report positive effects; while Cornwell et al. (2005), Angrist et al. (2009), Leuven et al. (2010), Patel and Rudd (2012) and Cohodes and Goodman (2014) do not find significant effects. Although there is little consensus on the reason behind the failure of many incentive programs, Dynarski (2008) notes that the incentives may be irrelevant for many students, and Angrist et al. (2014) report that one-third of the students in their study failed to fully understand a relatively simple grade-based incentive scheme. In other experiments on college students, academic support services have been combined with financial incentives. Results on the performance effects of these interventions are again mixed: Angrist et al. (2009) and Barrow et al. (2014) report strong effects; Angrist et al. (2014) find weak effects; and Miller et al. (2011) find no significant effects. Financial incentives are also controversial due to concerns that they might crowd out intrinsic incentives to study (see, e.g., Cameron and Pierce, 1994, and Gneezy et al., 2011). See Lavecchia et al. (2016) for a survey of financial incentives in higher education.

Web Appendix V.2 Details referred to in footnote 2

Commitment devices include purchase-quantity rationing of vice goods (Wertenbroch, 1998), deadlines (Ariely and Wertenbroch, 2002), commitments to future savings (Thaler and Benartzi, 2004), long-term gym membership contracts (DellaVigna and Malmendier, 2006), restricted access savings accounts (Ashraf et al., 2006), penalties for failing to hit work targets (Kaur et al., 2015) and Internet blockers (Patterson, 2016), while Augenblick et al. (2015) show that experimental subjects in the laboratory who are more present biased in the domain of work effort are more likely to use a commitment device. See Bryan et al. (2010) for a survey.

Web Appendix V.3 Details referred to in footnote 3

A small and recent literature in economics suggests that goal setting can influence behavior in other settings. Harding and Hsiaw (2014) find that goal setting can influence consumption: energy savings goals reduced energy consumption. Choi et al. (2016) find that goal setting can affect savings: goal-based cues increased savings into 401k accounts. Finally, goals can increase worker performance even in the absence of monetary incentives for achieving the goal: see Corgnet et al. (2015, 2016) for laboratory evidence (although Akın and Karagözoğlu, forthcoming, find no effect of goals), Goerg and Kube (2012) for field evidence and Goerg (2015) for a concise survey.

Web Appendix V.4 Details referred to in footnote 8

Morgan (1987) ran an experiment using one-hundred and eighty college students split into two control groups and three treatment groups. One treatment group set themselves goals for study time, pages to read and topics to cover. The second treatment group self-monitored but did not set goals. The third treatment group did both. The study tracked performance but did not report task completion. Average performance of the treated subjects was higher than that of control subjects. However, the paper does not report a statistical test of the performance difference between subjects who set goals and subjects in the controls.

Web Appendix V.5 Details referred to in footnote 9

A literature in psychology uses small-scale experiments to look at the effects of teacher-set goals on the learning of grade-school-aged children (e.g., LaPorte and Nath, 1976, Schunk, 1983, Schunk, 1984, Schunk and Rice, 1991, Schunk and Swartz, 1993, Schunk, 1996, and Griffee and Templin, 1997). There are important differences between teacher-set goals for grade-school-aged children and self-set goals for college students: first, college students can use self-set goals to regulate optimally their own behavior given their private information about the extent of their self-control problem; and second, in the school environment children are closely monitored by teachers and parents, which gives extrinsic motivation to reach goals (for instance, children might worry about explicit and implicit penalties, monetary or otherwise, for failing to achieve the goal set for them). Using a sample of eighty-four fourth-grade children, Shih and Alexander (2000) explore experimentally the effects of self-set goals (in particular, they study the effects of self-set goals for the number of fractions to solve in class on the ability to solve fractions in a later test).

Web Appendix V.6 Details referred to at the end of Section 4.3.3

In the Introduction we referred to evidence from educational environments that females have more self-control than males (e.g., Duckworth and Seligman, 2005, Buechel et al., 2014, and Duckworth et al., 2015). Consistent with gender differences in self-control, incentivized experiments suggest that men may be more present biased than women. When the earliest payment is immediate (no 'front-end delay'), which provides a test of quasi-hyperbolic discounting, McLeish and Oxoby (2007), Meier and Sprenger (2010) and Prince and Shawhan (2011) find that men are more present biased than women, while Tanaka et al. (2010) find no gender difference in present bias for rural Vietnamese (in Meier and Sprenger, 2010, when controls that include endogenous behavioral measures are included men remain more present biased than women, but the effect is no longer statistically significant). When rural Malawians are given an unexpected opportunity to reverse an earlier commitment, Giné et al. (forthcoming) find that men are more likely to reverse their earlier choice and instead choose an earlier but smaller payment. When the earliest payment is not immediate ('front-end delay'), no gender differences have been found (see Bauer and Chytilová, 2013, for rural Indians and Harrison et al., 2005, where the earliest payment is delayed by a month).

References

Akın, Zafer and **Karagözoğlu, Emin** (forthcoming). The role of goals and feedback in incentivizing performance. *Managerial and Decision Economics*

Angrist, Joshua, **Lang, Daniel**, and **Oreopoulos, Philip** (2009). Incentives and services for college achievement: Evidence from a randomized trial. *American Economic Journal: Applied Economics*, 1(1): 136–163

Angrist, Joshua, **Oreopoulos, Philip**, and **Williams, Tyler** (2014). When opportunity knocks, who answers? New evidence on college achievement awards. *Journal of Human Resources*, 49(3): 572–610

Ariely, Dan and **Wertenbroch, Klaus** (2002). Procrastination, deadlines, and performance: Self-control by precommitment. *Psychological Science*, 13(3): 219–224

Ashraf, Nava, **Karlan, Dean**, and **Yin, Wesley** (2006). Tying Odysseus to the mast: Evidence from a commitment savings product in the Philippines. *Quarterly Journal of Economics*, 121(2): 635–672

Augenblick, Ned, **Niederle, Muriel**, and **Sprenger, Charles** (2015). Working over time: Dynamic inconsistency in real effort tasks. *Quarterly Journal of Economics*, 130(3): 1067–1115

Barrow, Lisa, **Richburg-Hayes, Lashawn**, **Rouse, Cecilia Elena**, and **Brock, Thomas** (2014). Paying for performance: The education impacts of a community college scholarship program for low-income adults. *Journal of Labor Economics*, 32(3): 563–599

Bauer, Michal and **Chytilová, Julie** (2013). Women, children and patience: Experimental evidence from Indian villages. *Review of Development Economics*, 17(4): 662–675

Bell, David. E. (1985). Disappointment in decision making under uncertainty. *Operations Research*, 33(1): 1–27

Brunnermeier, Markus K and **Parker, Jonathan A** (2005). Optimal expectations. *American Economic Review*, 95(4): 1092–1118

Bryan, Gharad, **Karlan, Dean**, and **Nelson, Scott** (2010). Commitment devices. *Annual Review of Economics*, 2(1): 671–698

Buechel, Berno, **Mechtenberg, Lydia**, and **Petersen, Julia** (2014). Peer effects and students' self-control. *SFB 649 Discussion Paper 2014-024, Humboldt University*

Camerer, Colin and **Lovallo, Dan** (1999). Overconfidence and excess entry: An experimental approach. *American Economic Review*, 89(1): 306–318

Cameron, Judy and **Pierce, W David** (1994). Reinforcement, reward, and intrinsic motivation: A meta-analysis. *Review of Educational research*, 64(3): 363–423

Castleman, Benjamin L (2014). The impact of partial and full merit scholarships on college entry and success: Evidence from the Florida Bright Futures Scholarship program. *EdPolicyWorks Working Paper, University of Virginia*

Cha, Paulette and **Patel, Reshma** (2010). Rewarding progress, reducing debt: Early results from Ohio's performance-based scholarship demonstration for low-income parents. *Technical Report, MDRC*

Choi, James J, **Haisley, Emily**, **Kurkoski, Jennifer**, and **Massey, Cade** (2016). Small cues change savings choices. *Mimeo, Yale University*

Cohodes, Sarah R and **Goodman, Joshua S** (2014). Merit aid, college quality, and college completion: Massachusetts' Adams scholarship as an in-kind subsidy. *American Economic Journal: Applied Economics*, 6(4): 251–285

Corgnet, Brice, **Gómez-Miñambres, Joaquín**, and **Hernán-Gonzalez, Roberto** (2015). Goal setting and monetary incentives: When large stakes are not enough. *Management Science*, 61(12): 2926–2944

Corgnet, Brice, **Gómez-Miñambres, Joaquín**, and **Hernán-Gonzalez, Roberto** (2016). Goal setting in the principal-agent model: Weak incentives for strong performance. *CEDEX Discussion Paper 2016-09, University of Nottingham*

Cornwell, Christopher M, **Lee, Kyung Hee**, and **Mustard, David B** (2005). Student responses to merit scholarship retention rules. *Journal of Human Resources*, 40(4): 895–917

De Paola, Maria, **Scoppa, Vincenzo**, and **Nisticò, Rosanna** (2012). Monetary incentives and student achievement in a depressed labor market: Results from a randomized experiment. *Journal of Human Capital*, 6(1): 56–85

DellaVigna, Stefano and **Malmendier, Ulrike** (2006). Paying not to go to the gym. *American Economic Review*, 96(3): 694–719

Dhami, Sanjit (2016). *The Foundations of Behavioral Economic Analysis*. Oxford: Oxford University Press

Diamond, Peter and **Kőszegi, Botond** (2003). Quasi-hyperbolic discounting and retirement. *Journal of Public Economics*, 87(9): 1839–1872

Duckworth, Angela L and **Seligman, Martin EP** (2005). Self-discipline outdoes IQ in predicting academic performance of adolescents. *Psychological Science*, 16(12): 939–944

Duckworth, Angela L, **Shulman, Elizabeth P**, **Mastronarde, Andrew J**, **Patrick, Sarah D**, **Zhang, Jinghui**, and **Druckman, Jeremy** (2015). Will not want: Self-control rather than motivation explains the female advantage in report card grades. *Learning and Individual Differences*, 39: 13–23

Dynarski, Susan (2008). Building the stock of college-educated labor. *Journal of Human Resources*, 43(3): 576–610

Fang, Hanming and **Silverman, Dan** (2009). Time-inconsistency and welfare program participation: Evidence from the NLSY. *International Economic Review*, 50(4): 1043–1077

Fryer, Roland G (2013). Information and student achievement: Evidence from a cellular phone experiment. *Working Paper 19113, NBER*

Genesove, David and **Mayer, Christopher** (2001). Loss aversion and seller behavior: Evidence from the housing market. *Quarterly Journal of Economics*, 116(4): 1233–1260

Gill, David and **Prowse, Victoria** (2012). A structural analysis of disappointment aversion in a real effort competition. *American Economic Review*, 102(1): 469–503

Gill, David and **Stone, Rebecca** (2010). Fairness and desert in tournaments. *Games and Economic Behavior*, 69(2): 346–364

Giné, Xavier, **Goldberg, Jessica**, **Silverman, Dan**, and **Yang, Dean** (forthcoming). Revising commitments: Field evidence on the adjustment of prior choices. *Economic Journal*

Gneezy, Uri, **Meier, Stephan**, and **Rey-Biel, Pedro** (2011). When and why incentives (don't) work to modify behavior. *Journal of Economic Perspectives*, 25(4): 191–209

Goerg, Sebastian J (2015). Goal setting and worker motivation. *IZA World of Labor*, 178: 1–10

Goerg, Sebastian J and **Kube, Sebastian** (2012). Goals (th) at work–goals, monetary incentives, and workers' performance. *Preprint 2012/19, Max Planck Institute for Research on Collective Goods*

Gossner, Olivier and **Steiner, Jakub** (2016). Optimal illusion of control and related perception biases. *Discussion Paper 11478, CEPR*

Griffee, Dale T and **Templin, Stephen A** (1997). Goal setting affects task performance. In *Proceedings of the JALT 1997 International Conference on Language Teaching and Learning*, 21–26

Gruber, Jonathan and **Kőszegi, Botond** (2001). Is addiction rational? Theory and evidence. *Quarterly Journal of Economics*, 116(4): 1261–1303

Harding, Matthew and **Hsiaw, Alice** (2014). Goal setting and energy conservation. *Journal of Economic Behavior & Organization*, 107: 209–227

Harrison, Glenn W, **Lau, Morten I**, and **Rutström, Elisabet** (2005). Dynamic consistency in Denmark: A longitudinal field experiment. *Economics Working Paper 05-02, University of Central Florida*

Heath, Chip, **Larrick, Richard P**, and **Wu, George** (1999). Goals as reference points. *Cognitive Psychology*, 38(1): 79–109

Henry, Gary, **Rubenstein, Ross**, and **Bugler, Daniel** (2004). Is HOPE enough? Impacts of receiving and losing merit-based financial aid. *Educational Policy*, 18(5): 686–709

Hsiaw, Alice (2013). Goal-setting and self-control. *Journal of Economic Theory*, 148(2): 601–626

Hsiaw, Alice (2016). Goal bracketing and self-control. *Mimeo, Brandeis University*

Jain, Sanjay (2009). Self-control and optimal goals: A theoretical analysis. *Marketing Science*, 28(6): 1027–1045

Johnson, Dominic DP and **Fowler, James H** (2011). The evolution of overconfidence. *Nature*, 477(7364): 317–320

Kahneman, Daniel, **Knetsch, Jack L**, and **Thaler, Richard H** (1990). Experimental tests of the endowment effect and the Coase theorem. *Journal of Political Economy*, 98(6): 1325–1348

Kahneman, Daniel and **Tversky, Amos** (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2): 263–291

Kaur, Supreet, **Kremer, Michael**, and **Mullainathan, Sendhil** (2015). Self-control at work. *Journal of Political Economy*, 123(6): 1227–1277

Kőszegi, Botond and **Rabin, Matthew** (2006). A model of reference-dependent preferences. *Quarterly Journal of Economics*, 121(4): 1133–1165

Khwaja, Ahmed, **Silverman, Dan**, and **Sloan, Frank** (2007). Time preference, time discounting, and smoking decisions. *Journal of Health Economics*, 26(5): 927–949

Koch, Alexander K and **Nafziger, Julia** (2011). Self-regulation through goal setting. *Scandinavian Journal of Economics*, 113(1): 212–227

Koch, Alexander K and **Nafziger, Julia** (2016). Goals and bracketing under mental accounting. *Journal of Economic Theory*, 162: 305–351

Laibson, David (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics*, 112(2): 443–477

LaPorte, Ronald E and **Nath, Raghu** (1976). Role of performance goals in prose learning. *Journal of Educational Psychology*, 68(3): 260–264

Lavecchia, Adam, **Liu, Heidi**, and **Oreopoulos, Philip** (2016). Behavioral economics of education: Progress and possibilities. In E. Hanushek, S. Machin, and L. Woessman, editors, *Handbook of the Economics of Education, Vol. 5*, 1–74. North Holland

Leuven, Edwin, **Oosterbeek, Hessel**, and **van der Klaauw, Bas** (2010). The effect of financial rewards on students' achievement: Evidence from a randomized experiment. *Journal of the European Economic Association*, 8(6): 1243–1265

McLeish, Kendra N and **Oxoby, Robert J** (2007). Gender, affect and intertemporal consistency: An experimental approach. *Discussion Paper 2663, IZA*

Meier, Stephan and **Sprenger, Charles** (2010). Present-biased preferences and credit card borrowing. *American Economic Journal: Applied Economics*, 2(1): 193–210

Miller, Cynthia, **Binder, Melissa**, **Harris, Vanessa**, and **Krause, Kate** (2011). Staying on track: Early findings from a performance-based scholarship program at the university of New Mexico. *Technical Report, MDRC*

Morgan, Mark (1987). Self-monitoring and goal setting in private study. *Contemporary Educational Psychology*, 12(1): 1–6

Park, Young Joon and **Santos-Pinto, Luís** (2010). Overconfidence in tournaments: Evidence from the field. *Theory and Decision*, 69(1): 143–166

Patel, Reshma and **Rudd, Timothy** (2012). Can scholarships alone help students succeed? Lessons from two New York City community colleges. *Technical Report, MDRC*

Patterson, Richard W (2016). Can behavioral tools improve online student outcomes? Experimental evidence from a massive open online course. *Mimeo, United States Military Academy at West Point*

Pollak, Robert A (1968). Consistent planning. *Review of Economic Studies*, 35(2): 201–208

Prince, Jeffrey T and **Shawhan, Daniel L** (2011). Is time inconsistency primarily a male problem? *Applied Economics Letters*, 18(6): 501–504

Rabin, Matthew (2000). Risk aversion and expected-utility theory: A calibration theorem. *Econometrica*, 68(5): 1281–1292

Romer, David (1993). Do students go to class? Should they? *Journal of Economic Perspectives*, 7(3): 167–174

Schunk, Dale H (1983). Developing children's self-efficacy and skills: The roles of social comparative information and goal setting. *Contemporary Educational Psychology*, 8(1): 76–86

Schunk, Dale H (1984). Enhancing self-efficacy and achievement through rewards and goals: Motivational and informational effects. *Journal of Educational Research*, 78(1): 29–34

Schunk, Dale H (1996). Goal and self-evaluative influences during children's cognitive skill learning. *American Educational Research Journal*, 33(2): 359–382

Schunk, Dale H and **Rice, Jo Mary** (1991). Learning goals and progress feedback during reading comprehension instruction. *Journal of Literacy Research*, 23(3): 351–364

Schunk, Dale H and **Swartz, Carl W** (1993). Goals and progress feedback: Effects on self-efficacy and writing achievement. *Contemporary Educational Psychology*, 18(3): 337–354

Scott-Clayton, Judith (2011). On money and motivation: A quasi-experimental analysis of financial incentives for college achievement. *Journal of Human Resources*, 46(3): 614–646

Shih, Shu-Shen and **Alexander, Joyce M** (2000). Interacting effects of goal setting and self- or other-referenced feedback on children's development of self-efficacy and cognitive skill within the Taiwanese classroom. *Journal of Educational Psychology*, 92(3): 536–543

Smithers, Samuel (2015). Goals, motivation and gender. *Economics Letters*, 131: 75–77

Strotz, Robert Henry (1956). Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies*, 23(3): 165–180

Suvorov, Anton and **Van de Ven, Jeroen** (2008). Goal setting as a self-regulation mechanism. *Working Paper 122, CEFIR NES*

Tanaka, Tomomi, **Camerer, Colin F**, and **Nguyen, Quang** (2010). Risk and time preferences: Linking experimental and household survey data from Vietnam. *American Economic Review*, 100(1): 557–571

Thaler, Richard and **Benartzi, Shlomo** (2004). Save more tomorrow^{TM}: Using behavioral economics to increase employee saving. *Journal of Political Economy*, 112(S1): S164–S187

Weinstein, Neil D (1980). Unrealistic optimism about future life events. *Journal of Personality and Social Psychology*, 39(5): 806–820

Wertenbroch, Klaus (1998). Consumption self-control by rationing purchase quantities of virtue and vice. *Marketing Science*, 17(4): 317–337

Wu, George, **Heath, Chip**, and **Larrick, Richard** (2008). A prospect theory model of goal behavior. *Mimeo, University of Chicago*