Improving College Instruction through Incentives

Andy Brownback
University of Arkansas

Sally Sadoff*
UC San Diego

March 27, 2018

Abstract

In a field experiment, we test the impact of performance-based incentives for community college instructors. We estimate that instructor incentives improve student performance on objective course exams by 0.2 standard deviations, increase course grades by 0.1 standard deviations, and reduce course dropout rates by 17%. The largest effects are among part-time adjunct instructors. To test for potential complementarities, we also examine the impact of instructor incentives in conjunction with student incentives and find no evidence that the incentives are more effective in combination. Our instructor incentives are framed as losses and distributed in the form of upfront bonuses that instructors pay back at the end of the semester if they do not meet performance targets. We elicit instructors' contract preferences and find that, at baseline, instructors prefer to work under gain-framed contracts with rewards distributed at the end of the semester. However, after experiencing the loss-framed incentives, instructors significantly increase their preferences for them.

JEL codes: I23, M52, M55, C93
Keywords: college instruction, incentives, loss-framed contract preferences

---

*Corresponding Author: Sally Sadoff, UC San Diego, Rady School of Management, 9500 Gilman Drive, La Jolla, CA 92093; sadoff@ucsd.edu. We gratefully acknowledge Ivy Tech Community College and our invaluable partners there, Ronald Sloan and Neil Anthony. Eurika Bennett, Jaclyn Fisher, Ryan Johnson, Jacklyn Redwine, Tammy Jo Robbins, and Alisa Wells provided truly outstanding assistance with data collection and experimental execution. The project was made possible by the generous financial support of the Laura and John Arnold Foundation and the Yankelovich Center for Social Science Research. This research was conducted with approval from the University of Arkansas and UC San Diego Institutional Review Boards.

1 Introduction

Over the last several decades, both the returns to higher education and postsecondary enrollment rates have increased, while completion rates have stagnated (Oreopoulos and Petronijevic, 2013). Attainment is particularly low at two year community colleges, which serve about forty percent of all undergraduates and half of those who eventually earn a four year degree (Shapiro, 2017; McFarland et al., 2017). A growing literature examines the impact of policy interventions aimed at improving postsecondary student performance, including lowering the costs of college attendance, providing students with information and support services, and offering students performance-based incentives (Deming and Dynarski, 2009; Lavecchia et al., 2014; Evans et al., 2017, provide reviews). These policies have generally targeted students, largely ignoring the role of college instructors.

At the same time, a growing body of research demonstrates the importance of postsecondary instruction. A one standard deviation (SD, hereafter) increase in instructor quality improves student performance by an estimated $0.05 - 0.3$ SD, with effects generally smaller at selective universities and larger at non-selective institutions similar to community colleges.[1] Recent work examines the extent to which postsecondary institutions adjust personnel policies – such as teaching allocations and salaries – in response to higher instructor performance and productivity (Courant and Turner, 2017; De Vlieger et al., 2017). But, to our knowledge, no prior study has explored whether it is possible to improve the effectiveness of college instructors.

The dearth of interventions targeting college instructors stands in sharp contrast to a large literature on improving teacher performance at the elementary and high school levels. This gap is particularly surprising given the scope for personnel policies in higher education, where teaching assignments and employment contracts are generally more flexible than in most K-12 settings. This flexibility has increased with the sharp rise in part-time adjunct instructors who work under short term contracts and teach courses at lower cost than full-time faculty (Ehrenberg, 2012; McFarland et al., 2017).[2]

We fill the gap in the literature by experimentally testing the impact of performance-based incentives among community college instructors. Our primary question is whether incentives for instructors can improve postsecondary student performance. Secondarily, we test whether instructor incentives can be more effective in combination with

[1] See Hoffmann and Oreopoulos (2009); Carrell and West (2010); Braga et al. (2016); Brodaty and Gurgand (2016); Bettinger et al. (2014); De Vlieger et al. (2017).

[2] There is ongoing debate about the impact of this shift on student achievement (Ehrenberg and Zhang, 2005; Bettinger and Long, 2006; Hoffmann and Oreopoulos, 2009; Bettinger and Long, 2010; Figlio et al., 2015; Rogers, 2015; Ran and Xu, 2016).

incentives for students. Finally, in order to explore their feasibility as a personnel policy, we examine instructor preferences for the incentive contracts we offer.

Community colleges have successfully increased access to postsecondary education, particularly among underrepresented groups, but struggle with low student achievement (Bailey et al., 2015). Estimated six-year degree rates for students starting at community colleges remain below forty percent compared to above sixty percent at four year colleges (Shapiro, 2017). Low on-time degree rates stem in part from poor course performance. Comparing students in the same state college system, Ran and Xu (2016) estimate that only two-thirds of community college students earn a C-grade or better in introductory courses compared to about three-quarters of students at four year colleges. Course completion rates follow a similar pattern: about one in six community college students withdraws from introductory courses compared to about one in eleven students at four year colleges. Low course completion rates can also cause students to lose their financial aid, increasing the likelihood of dropping out with debt. We therefore focus on both course performance and course completion as key outcomes of interest.[3]

We conducted our experiment at several campuses of a statewide community college in the Fall and Spring semesters of the 2016-2017 school year. The study included sixteen different departments with over 6,000 student-course observations. In the Fall semester, we randomly assigned instructors to one of two treatment groups: Instructor Incentives or Control. In the Instructor Incentives group, instructors received performance bonuses of $50 per student who received a 70 percent or higher on an objective, externally-designed course exam ("passed," hereafter). We framed the incentives as *losses* – i.e., bonuses instructors would lose if they did not meet performance targets. To implement the loss framing, we gave instructors upfront bonuses at the *beginning* of the semester equivalent to the amount they would receive if half of their students passed the exam. At the end of the semester, if fewer than half of an instructor's students passed the exam, the instructor returned the difference between their final reward and the upfront bonus. If over half of the students passed the exam, the instructor received additional rewards.

[3]Course withdrawals – dropping a course after two or more weeks – provide neither credit nor refund. Federal financial aid requires students to maintain a 67% completion rate. Four community colleges (including our partner institution) are among the top five public postsecondary institutions producing the highest ratio of dropouts with debt to graduates. See J. Barshay, "3.9 Million Students Dropped Out of College with Debt in 2015 and 2016," *U.S. News*, (2017; `https://www.usnews.com/news/data-mine/articles/2017-11-07/federal-data-show-39-million-students-dropped-out-of-college-with-debt-in-2015-and-2016`)

In the Spring semester, we introduced Combined Incentives, which offered incentives to students in conjunction with incentives to instructors. Incentivized students received free tuition for one summer course (worth approximately $400) if they passed the exam. We assigned student incentives at the section level, cross-randomizing them with the existing assignment of instructor incentives. This yields four treatment groups: Instructor Incentives *only*, Student Incentives *only*, Combined Incentives (instructor incentives and student incentives), and Control. In order to explore potential complementarities between instructor and student incentives, we examine whether Combined Incentives are more effective than Instructor Incentives alone.

Finally, we used incentive compatible mechanisms to elicit instructors' contract preferences both at baseline when they enrolled in the study and at the end of the semester after incentivized instructors had experienced the contracts. We compare the the loss-framed contract to a more standard gain-framed contract in which rewards are distributed at the end of the semester ("loss" and "gain" contracts, respectively). Our elicitation captures the difference in the per-student incentive amount that would make an instructor indifferent between working under the loss versus the gain contract.

We find that Instructor Incentives increase exam performance by an estimated 0.20 SD ($p \approx 0.000$). The impact carries over to overall course performance, where course grades improve by 0.11 SD ($p = 0.039$). These effects operate at both the extensive and intensive margins. Instructor Incentives increase course completion by 3.7 percentage points ($p = 0.050$), which represents a 17% reduction in the course withdrawal rate; and also increase exam scores conditional on completion by 0.082 SD ($p = 0.040$). At the instructor level, the effects of incentives are largest among adjunct instructors (0.26 SD, $p \approx 0.000$) with smaller effects among full-time faculty (0.13 SD, $p = 0.120$). We find no evidence that Combined Incentives are more effective than Instructor Incentives.[4] Rather than finding evidence of complementarities between instructor incentives and student incentives, we find suggestive evidence that Combined Incentives are *less* effective than Instructor Incentives alone.

The incentives we test are low-cost, short-term and targeted at the individual course level. Our results suggest, however, that the effects we find could scale up. First, the effects are consistently positive across a wide range of courses. Second, the effects are sustained across multiple semesters with larger estimated effects in the second semester that incentives are offered. Third, there is no evidence that instructors or students substitute their effort away from unincentivized courses to incentivized courses. And the

[4] We also find little evidence that Student Incentives have meaningful effects when offered alone. We note that we did not power the experiment to separately estimate the impact of Student Incentives.

estimated treatment effects are sustained as instructors and students are exposed to more incentivized courses. Finally, the impact of incentives on course completion spills over to non-program courses. These courses experience similar declines in withdrawal rates to incentivized courses, suggesting that the incentives encourage general enrollment during the treatment semester. We do not, however, find any evidence of an impact on enrollment in subsequent semesters.

Turning to instructors' contract preferences, we find two striking results. First, at baseline, instructors significantly prefer gain contracts to loss contracts. They are willing to give up 9-12% of the $50 per-student incentive payment in order to work under a gain contract rather than a loss contract. Second, after one semester of working under loss contracts, incentivized instructors significantly increase their preferences for them. The effects are large enough that instructors become (close to) indifferent between loss and gain contracts.

To our knowledge, this is the first study to demonstrate that an intervention can improve instructor effectiveness at the post secondary level. And the first to examine performance-based incentives among college instructors. Our results contribute to a growing literature examining teacher incentives at the elementary and high school levels that has found mixed results. While non-experimental studies and experimental studies in developing countries have found that teacher incentives can improve performance, experimental studies in the U.S. have largely failed to demonstrate effectiveness.[5] We based the design of the incentives in our study on Fryer Jr et al. (2012, 2018), which is the only prior experimental study in the U.S. to find a positive impact of teacher incentives. They test upfront, loss-framed incentives among elementary and middle school teachers and estimate effects of 0.12 SD on math test scores pooling across two years of the experiment. Our findings that similarly structured incentives are effective among college instructors suggests that the impact of loss-framed incentives on teacher performance may replicate across contexts.[6]

We also add to a small set of existing studies that have found conflicting results when comparing incentives offered both alone and in combination.[7] In line with our

[5] Neal (2011); Fryer Jr et al. (2018) provide reviews. For experimental studies in developing countries, see Glewwe et al. (2010); Muralidharan and Sundararaman (2011); Duflo et al. (2012); Loyalka et al. (2016); Barrera-Osorio and Raju (2017). For experimental studies in the U.S., see Glazerman et al. (2009); Fryer (2013); Springer et al. (2011, 2012).

[6] Unlike Fryer Jr et al. (2012, 2018), we do not attempt to compare loss- and gain-framed contracts. They find that gain-framed contracts have an estimated impact of 0.05 SD pooling across two years of the experiment. Our incentives also differ. We base rewards on threshold achievement levels, while Fryer Jr et al. (2012, 2018) used the pay-for-percentile structure developed by Barlevy and Neal (2012).

[7] A large literature examines student incentives alone and generally finds small effects (see reviews by Sadoff, 2014; Fryer, 2017). In a community college context, Barrow et al. (2014); Patel et al. (2013)

results against complementarities between instructor and student incentives, List et al. (2012) find little evidence of complementarities between incentives for students, parents, and tutors in an experiment in U.S. elementary schools. In contrast, Behrman et al. (2015) find that incentives for teachers and students in Mexican high schools were more effective when offered in combination than when offered separately.[8] The differing results across studies could be driven by differences in complementarities between instructor and student effort in the production function; or could also be due to differences in the strategic response of instructors and students to each others' effort choices (see De Fraja et al., 2010, for discussion).

Finally, our study contributes to personnel economics by examining employee preferences for loss-framed contracts in a high-stakes, natural environment using employee salaries. The motivational power of loss contracts is consistent with a large literature in behavioral economics demonstrating *loss aversion*, under which having to pay back (or "lose") part of the bonus is more painful than failing to earn (or "gain") the equivalent amount at the end of the semester (Kahneman and Tversky, 1979). A growing body of laboratory and field studies demonstrates that framing incentives as losses can increase worker effort compared to more traditional gain-framed incentives.[9]

Despite their potential impact on productivity, however, explicit loss-framed contracts are not widely prevalent. One reason for this may be that workers find the loss-framed contracts aversive and prefer to work under gain-framed contracts. If this is the case, employers may need to increase employee compensation in order to retain employees who work under loss-framed contracts, offsetting the improved productivity. While standard behavioral models predict that workers will prefer gain-framed contracts, the limited empirical evidence from laboratory and online studies finds a preference for loss-framed contracts (Imas et al., 2016; De Quidt, 2017).[10] Our study

find that performance-based scholarships for students modestly improve GPA, credit accumulation, and degree attainment. In contrast to our low-cost, short-term incentives, these scholarships were expensive ($1000-$4000) and long-term.

[8] We note that the combined intervention in Behrman et al. (2015) had programmatic elements that were not included in the individual interventions. Using observational data, Geng (2017) finds evidence of complementarities between a grade retention policy incentivizing students and an accountability scheme incentivizing teachers and schools. In related work, Mbiti et al. (2016) find complementarities between teacher incentives and school resources.

[9] See Brooks et al. (2012); Hossain and List (2012); Fryer Jr et al. (2012, 2018); Hong et al. (2015); Armantier and Boly (2015); Imas et al. (2016). In online studies, DellaVigna and Pope (2016) and De Quidt et al. (2017) do not find significant differences between loss- and gain-framed incentives. Studies comparing loss- and gain-framed incentives outside of work settings find mixed results (e.g., List and Samek, 2015; Levitt et al., 2016; Englmaier et al., 2018).

[10] Models using the status quo as the reference point (e.g., Tversky and Kahneman, 1991) predict that individuals will work harder under loss contracts conditional on the endowment (i.e., the upfront bonus) being incorporated as the status quo. If the distribution of possible outcomes (i.e., final

is the first to examine contract preferences both before and after working under loss-framed incentives. Interestingly for both theory and policy, our results suggest that people who experience loss-framed contracts do not judge those experiences as negatively ex-post as they did ex-ante.[11]

Our study demonstrates that incentives can substantially improve postsecondary instruction at low cost. The effects of our incentives are similar in size to improving teacher quality by a standard deviation and have an expected cost of about $25 per student-course. Our community college context may make incentives particularly powerful, since low cost rewards can provide a substantial bonus relative to baseline pay. This is particularly the case for adjunct instructors, for whom the expected incentive was equivalent to approximately twenty percent of their salary.[12] The dramatic impact of our incentives on adjunct instructors suggests that there could be substantial gains from reconsidering the contracts offered to part-time instructors. These changes could not only significantly improve student outcomes but are also feasible from a policy perspective given the preferences of instructors, the low cost of the incentives, and the short-term contracting used to hire adjunct faculty.

In the remainder of the paper, Section 2 describes the experimental design, Section 3 presents the results for the impact of incentives on performance, Section 4 examines instructors' contract preferences, and Section 5 concludes.

2 Experimental design

2.1 Setting and recruitment

We conducted the experiment in the 2016-2017 school year at Ivy Tech Community College of Indiana (Ivy Tech). Ivy Tech is Indiana's largest public postsecondary institution and the nation's largest singly accredited, statewide community college system serving nearly 170,000 students annually. Our sample includes courses from several campuses in the East Central and Richmond regions: Anderson, Connorsville, Marion, Muncie, New Castle, and Richmond.

The East-Central and Richmond regions respectively serve communities in the 4th

rewards) is taken as the reference point (e.g., Kőszegi and Rabin, 2006), then the contract framing should be irrelevant for both effort and preference. Thus, our "standard" behavioral model refers to one assuming the status quo as the reference point. See Imas et al. (2016) for discussion of the theory.

[11]This finding is in line with Kermer et al. (2006), which argues that the affective experience of losses is less painful than people expect it to be.

[12]For comparison, Fryer Jr et al. (2012, 2018) offered expected bonuses equivalent to about eight percent of the average teacher's salary.

and 8th percentile of national median income. Over 60% of their student body is eligible for Pell Grants, placing them in the 90th percentile for community colleges. Their fall-to-fall retention rates of full-time students hover around 40%, just above the bottom 10% of community colleges. Overall, only 24% of their full-time, first-time students will graduate or transfer to a four-year institution within 3 years, also just above the bottom 10% of community colleges (NCCBP, 2014).

Our study includes a broad range of departments: Accounting, Anatomy and Physiology, Art History, Biology, Business Operations Applications and Technology, Business, Communications, Criminology, English, Health Sciences, Math, Nursing, Psychology, Physical Science, Software Development, and Sociology. We determined course eligibility based on whether the course included an objective course-wide exam. The exams were developed at the department level as part of a region-wide effort to standardize measures of learning across campuses and instructors. To ensure instructors were not able to "game" our incentives, department heads agreed to maintain confidentiality of the exam prior to its administration. Any instructor who taught at least one section of an eligible course was invited to participate.

In the Fall 2016 semester, Ivy Tech identified approximately 150 eligible instructors. Ivy Tech administrators recruited these instructors by email and in person. We then enrolled interested instructors in the study through an online survey. The enrollment period began August 15th, 2016 and ended September 6th, 2016 with a final total of 108 enrolled instructors, 90% of our recruitment goal of 120 and 72% of all eligible instructors. The randomization (detailed in Section 2.4) was based on the students enrolled in a given course as of the Ivy Tech census date, September 2nd, 2016. The census date is at the end of the second week of courses and is the final date students can receive a refund for a course. By delaying the randomization, we can control for selective take-up or attrition resulting from treatment assignment. Additionally, we can ensure that our estimates of withdrawals are not influenced by the natural add and drop cycles during the first two weeks of class.[13]

Fall instructors teaching eligible courses in the Spring 2017 semester were automatically re-enrolled. Of the 108 participating instructors in the fall, 74 were eligible in the spring and all but one elected to continue participation (as discussed in Section 2.4, there were no differences in eligibility by treatment group). We also recruited new instructors. The recruitment followed the same procedure as in the Fall 2016 semester, with Ivy Tech administrators emailing 74 eligible instructors who either had chosen

[13]Adding a course after the census date is rare and requires special permission from the instructor of record and the regional academic officer.

not to participate in the fall semester or were newly eligible. The enrollment period began January 20th, 2017 and ended February 3rd, 2017. An additional 26 instructors signed up, bringing the spring semester total to 99 participating instructors. Including continuing instructors from the fall, 66% of eligible instructors participated in the spring. As in the fall, the spring randomization was based on enrollment as of the spring semester census date, January 30th, 2017. Over both semesters, 134 instructors participated in the study, 93% of our recruitment goal of 144.

2.2 Treatments

We test two cross-cutting incentive schemes: incentives for instructors and incentives for students, which yields four treatments: Instructor Incentives only, Student Incentives only, Combined Incentives and Control. In the Instructor Incentives and Combined Incentives treatments, instructors received $50 per student who scored a 70% or higher on the objective course exam ("passed," hereafter).[14] Instructors received incentives for all students in all of their eligible sections. At the beginning of the semester, the University of Arkansas distributed checks for upfront bonuses equivalent to the amount instructors would earn if 50% of their students passed the exam. For example, an instructor who taught one section with 20 students would receive an upfront check for $500. At the end of the semester, if fewer than 50% of the students passed the exam, the instructor was responsible for returning the difference between the final bonus amount and the upfront bonus. If more than 50% of the students passed, the instructor received an additional bonus check.[15]

At the beginning of the semester, we notified instructors of their treatment assignment. We emailed instructors assigned to Instructor or Combined Incentives a list of their incentivized sections and an estimate of the upfront incentive payment they would receive. In order to clarify details and give instructors a chance to fill out the accounting forms in person, we held information sessions on each of the four primary campuses (Anderson, Marion, Muncie, and Richmond). One information session each semester was broadcast online for those who could not attend in person. Instructors who did not attend the session could sign the forms through our Ivy Tech partners or electronically. The upfront bonus payment was issued once the forms were signed.

[14] The health sciences and nursing courses had thresholds for passing that exceeded 70%. In these cases, we considered "passing" to be the pre-existing requirement of 75%.

[15] Recent work demonstrates that under a prospect theory model with both loss aversion and diminishing sensitivity (i.e., utility is convex in losses and concave in gains), contracts like ours – offering both bonuses and penalties for performance above and below a threshold respectively – can increase worker effort compared to pure bonus or pure penalty contracts (Armantier and Boly, 2015).

The average upfront bonus was $726 and the average final bonus was $662. Fifty-five percent of instructors owed money back at the end of the semester with an average repayment of $308. We had high rates of compliance for the upfront bonuses – 98% of instructors in the fall and 94% of instructors in the spring complied with the upfront payments.[16] Compliance with repayments varied across the two semesters – in the fall, 93% of instructors who owed money complied with repayment (96% of money owed), and in the spring 78% of instructors who owed money complied with repayment (83% of money owed). The lower repayment rate in the spring may have been due to instructors knowing that the study would not continue after the spring semester, a concern that would not be present if this were a system-wide policy. If instructors expected not to make repayments, this would likely lower the impact of incentives and thus our ability to detect treatment effects.

In the Student Incentives and Combined Incentives treatments, students received free tuition for one summer course if they passed their exam in the treated course. We designed the incentives in partnership with Ivy Tech to satisfy several administrative constraints. Offering cash incentives is not feasible, as cash rewards can crowd out existing financial aid for certain students. Relatedly, because summer enrollment may not be covered by Pell Grants, summer scholarships can help lower a student's debt burden beyond what a fall or spring scholarship could do. The summer scholarship incentives are also attractive from a cost perspective. A summer scholarship has a face value of $400 but has an expected marginal cost of only about $97 given realized pass rates of 44.7% and take up rates of 54.4% in sections offering student incentives.

In Spring 2017, we informed instructors which (if any) of their sections would receive incentives for students and the basic design of the incentives. An Ivy Tech administrator described the incentives to students (in person for traditional classes and through a video for online classes). Participating students received a refrigerator magnet reminding them of the details (Appendix Figure A.1). Students enrolled in the program by signing up in their class or through an online survey. Of the 1035 students offered incentives, 772 (74.6%) actively consented to participate and 48 (4.6%) actively declined to participate. Our primary analysis is at the intent-to-treat level and does not depend on whether a student chose to participate in the program.[17]

[16]The remainder did not fill out the paperwork to receive payments (3 instructors) or did not cash the upfront payment check (1 instructor).

[17]Consent did not affect our access to anonymous student-level data but did affect whether or not we could distribute summer tuition vouchers to students.

2.3 Survey

All participating instructors filled out a short online survey in order to enroll in the experiment. We asked instructors participating in the fall semester to fill out a mid-year survey at the end of the fall semester before they learned their final payment in December, 2016 (instructors new to the program in the spring filled out the enrollment survey at this time). We asked instructors participating in the spring semester to fill out a year-end survey in May, 2017. Response rates were 87% for the mid-year survey (96% in Instructor Incentives and 77% in Instructor Control). Response rates were lower for the year-end survey: 67% (83% in Instructor Incentives and 49% in Instructor Control).

In the enrollment and mid-year surveys, we elicited instructors' preferences for loss-framed relative to gain-framed contracts (see Appendix B for preference elicitation questions). First, we asked instructors to choose which contract they would prefer to receive if both contracts paid $50 per student. Then, we used a multiple price list in which instructors made a series of decisions between an "Advance Bonus"– a loss-framed contract that provided half of the total possible reward upfront – and an "End Bonus" – a gain-framed contract that paid all rewards at the end of the semester. Our multiple price list elicited preferences between the loss-framed contract with a bonus of $50 per student and 13 different gain-framed contracts with a bonus of $X per student with $X \in \{60, 55, 54, 53, 52, 51, 50, 49, 48, 47, 46, 45, 40\}$. In order to ensure that the surveys were incentive-compatible, we randomly selected one choice from one respondent at random to determine that respondent's contract.[18]

Contract preferences may be confounded by time preferences because more impatient instructors may express a relatively stronger preference for loss-framing due to the earlier arrival of the payments (and vice versa for more patient instructors and gain-framed contracts). In order to separately identify contract preferences from time preferences, we also elicited instructors' preferences over receiving unconditional cash awards at the beginning versus the end of the semester. Similar to the multiple price list for contracts, instructors made thirteen decisions between a $500 bonus at the beginning of the semester and a bonus of $B at the end of the semester with $B \in \{600, 550, 540, 530, 520, 510, 500, 490, 480, 470, 460, 450, 400\}$.

In all surveys, we asked instructors about their subjective well being and attitudes towards teaching. In the mid-year and year-end surveys we also asked about their time use and personal expenditures on instruction. For instructors in the treatment group,

[18] In the randomly-chosen decisions from both the fall and the spring, the instructor selected the loss contract and so received the same incentives as the other treatment instructors.

we additionally asked how they had used their upfront payments and their expectations about their final reward (e.g., whether they expected to receive additional rewards or owe money back).

2.4 Randomization

We first describe the randomization of instructors to receive incentives in the fall and spring semesters (See Appendix Figure A.2 for a summary). We then describe the randomization of individual sections to receive incentives for students in the spring semester. The randomization and analysis follows our pre-analysis plan.[19]

In the Fall 2016 semester, we assigned instructors to either Instructor Incentives or Instructor Control. We used a block randomized design, stratifying our instructors by department and instructor type (adjunct or full-time faculty).[20] We intended to stratify at a finer level, but that would have undermined the random assignment. To ensure balance between treatment and control, we tested for significant differences in course-level characteristics: courses per instructor, students enrolled per instructor, the percentage of courses with a corresponding remedial section (Co-Req), as well as instructors' time preferences and instructors' contract preferences elicited in the enrollment survey. We also tested for significant differences in student-level characteristics: gender, age, race, accumulated credit hours, and Grade Point Average (GPA). For each characteristic, we specified that we would re-randomize in the event that differences were significant with a p-value < 0.15.[21]

In Spring 2017, we conducted the randomization in two stages. First, we determined if an instructor would receive incentives. Next, we assigned which sections would receive student incentives.

For the instructor incentive stage of the randomization, we independently assigned continuing instructors who participated in the fall, and instructors who were new to the program. Of 55 instructors assigned to Instructor Incentives in the fall, 37 taught eligible courses in the spring. Of the 53 instructors assigned to Instructor Control in the fall, 37 taught eligible courses in the spring. Continuing instructors were assigned to the same treatment they received in the fall. The exceptions to this are: (1) One eligible instructor assigned to Instructor Incentives in the fall opted out of the spring

[19]We pre-registered our analysis plan. See https://osf.io/fbxpw/. We note below deviations from the pre-analysis plan due to data or experimental constraints.

[20]For some departments, it was impossible to stratify on both instructor type and department. In these cases, we pooled courses across departments and stratified on instructor type.

[21]We used a probit regression and regressed the treatment assignment on the characteristics. We re-randomized if any coefficients were significant with $p < 0.15$.

semester of the study; and (2) two instructors assigned to Instructor Control in the fall received Instructor Incentives in the spring. In order to encourage survey completion and continued participation among Control instructors, we told them they would have a chance to receive incentives in the spring. Accordingly, we randomly re-assigned two instructors. Therefore, we have 38 continuing instructors assigned to spring Instructor Incentives and 35 continuing instructors assigned to spring Instructor Control. Continuing fall instructors were well-balanced on baseline characteristics. The only significant difference is that instructors continuing in Instructor Incentives taught 0.53 more sections on average than those continuing in Instructor Control, significant at the $p < 0.05$ level.

New instructors were assigned to Instructor Incentives or Instructor Control following the same procedure as in the fall randomization. While we checked the balance of these characteristics among the full sample of instructors, we ran the randomization for new instructors independently to ensure that new spring instructors underwent the exact same assignment process as new fall instructors.

For the student incentive stage of the randomization, we assigned sections to receive student incentives within each of the instructor incentive assignments. For instructors assigned to receive incentives, we selected half of all their sections to receive Student Incentives (making them "Combined Incentives" sections), while the other half remained as Instructor Incentives only sections. In order to maximize within-instructor variation, any instructor with multiple sections had half of their sections assigned to receive Combined Incentives. Instructors with an odd number of sections were randomly rounded up or down. For instructors who taught one section, half were assigned to receive Combined Incentives and half received Instructor Incentives only.

For instructors assigned to Instructor Control, we first randomized half the instructors to a pure control group (no instructor or student incentives). Among the other half of Instructor Control instructors, we selected half of their sections to receive Student Incentives only and the other half of their sections to remain as pure control, following the same procedure described for the Instructor Incentive group. This asymmetrical method of assigning Student Incentives to instructors based on their instructor incentive assignment preserves a pure control (no student or instructor incentives) group. It also allows for a more powerful within-instructor test of complementarity between Instructor Incentives and Student Incentives. We balanced the student incentives randomization on all of the same characteristics as the instructor incentive assignment.

2.5 Analysis

We test two hypotheses: first, that instructor incentives improve student outcomes; and second, that instructor incentives have larger effects in combination with student incentives than they do alone. We estimate the following equation using a random effects linear regression model with standard errors clustered at the instructor level:[22]

$$Y_{i,j,s} = \beta_0 + \beta_1 Z^1_{s,j} + \beta_2 Z^2_{s,j} + \beta_3 Z^3_{s,j} + \beta_4 X_i + \beta_5 X_s + \beta_6 X_j + U_j + \epsilon_{i,s}$$

where $Y_{i,s,j}$ is the outcome for student i in section s taught by instructor j; $Z^t_{s,j}$ is an indicator variable for whether section s taught by instructor j is assigned to treatment $t = \{1, 2, 3\}$ with 1 = Instructor Incentives, 2 = Student Incentives, and 3 = Combined Incentives; X_i represents a vector of student covariates (age, race, gender, baseline credits); X_s represents course-specific covariates (academic department, and whether it is a co-requisite course); X_j represents instructor-specific covariates (full-time or adjunct, time preference, and contract preference); U_j represents the instructor-specific random effect; and $\epsilon_{i,s}$ is the error term which, due to the randomization, is mechanically uncorrelated to the $Z^t_{s,j}$ terms.

Since we partition our sections into the three treatments or control, β_1, β_2, and β_3 measure the full effects (rather than marginal effects) of Instructor Incentives alone, Student Incentives alone, and Combined Incentives, respectively. Based on our realized sample size and randomization, we estimate a minimum detectable effect size (MDES) for our primary outcome – performance on the objective exam – of 0.17 SD for Instructor Incentives, identical to our pre-analysis plan; and of just under 0.22 SD for Combined Incentives, compared to a MDES of 0.2 SD in our pre-analysis plan. We powered the study with a larger MDES for Combined Incentives given their higher cost and our interest in testing the hypothesis that Combined Incentives have larger effects than Instructor Incentives alone. The pairwise test, $\beta_3 > \beta_1$, itself has a MDES of 0.25 SD. We did not have a large enough sample size to adequately power a test of Student Incentives alone, β_2, or the full test of complementarities between istructor and student incentives, which compares the effect of Combined Incentives to the sum of the effect of Instructor Incentives and the effect of Student Incentives, $\beta_3 > \beta_1 + \beta_2$.

[22]For binary outcomes, we use a random effects probit model with clustered standard errors.

3 Results

We collected course data for 6,241 student-course observations in 383 sections. Our administrative dataset does not have demographic characteristics for 175 student-course observations, leaving us with a final sample size of 6,066 student-course observations (3,575 unique students). We are missing exam data from eight instructors in the fall semester and three instructors in the spring semester, yielding 5,839 student-course observations with valid exam data.[23] There are no differences by treatment in the rate of missing baseline characteristics or exam data (Appendix Table A.2). Nonetheless, to address concerns about missing data, we run our analysis on the exam data and the course data separately.

Table 1 reports means and proportions (with standard errors clustered by instructor) for baseline characteristics by semester and treatment for the following student-level characteristics: age, gender, race/ethnicity, total credits accumulated at baseline, baseline GPA, and whether GPA is missing (all newly entering students and some students returning after long absences have missing baseline GPAs); instructor-level characteristics: full-time or adjunct, total sections in the study, students per section, and elicited contract and time preferences; and section-level characteristics: whether the course section is a co-requisite (the Ivy Tech co-requisite course model is a form of remedial education for under-prepared students that operates concurrently with the enrolled course).[24]

To show we are balanced on baseline characteristics, we report the p-value from an F-test of equality across all groups within each semester. For the spring semester, we also report significant differences of means from binary tests comparing each treatment group to the control group. Of the forty-eight pairwise tests of differences we conduct,

[23] Appendix Table A.1 shows the distribution of students, instructors, and courses with valid exam data and with course data across the the two semesters by treatment. Exam scores are not individually recorded in the administrative data and had to be collected by our data collection team for the study. One instructor left Ivy Tech in the middle of the fall semester and the replacement instructor did not submit exams to our data collection team. Also in the fall semester, six instructors (teaching seven courses) in the Business, Operations, Applications, and Technology department recorded grades for their exams as pass or fail instead of recording scores. One additional instructor from the fall and three from the spring failed to submit their exams to our data collection team for unknown reasons and were unavailable when we attempted to follow-up.

[24] Contract and time preferences are included in our analysis as an indicator variables for above or below the median preference for loss-contract framing (relative to gain-contract framing) and end of semester payments (relative to start of semester payments), respectively. Indicator variables avoid the need to assign values to top- and bottom-coded data. If we cannot estimate an instructor's preference in the fall or spring semester due to missing or incomplete surveys, we substitute the value measured in the other semester. This affects contract preference values for 2 fall and 11 spring instructors and time preference values for 2 fall and 10 spring instructors.

three are significant at the 10% level, slightly less than what would be expected by chance (none are significant at the 5% level).

3.1 Effects of incentives on exam scores

We first examine the effect of treatment on the directly incentivized outcome: performance on the objective course exam. Unless otherwise noted, test scores are normalized within department to have a mean of zero and a standard deviation of one. Students who withdrew from the course after the drop deadline are coded as having received a zero on the final exam.

Table 2 displays the results of our random effects linear regression analysis for the full year (columns 1-3) and by semester (columns 4 and 5 for fall and spring, respectively).[25] We also report the p-value from a test of equality of the effects of Instructor Incentives and Combined Incentives. In column 1, we include only indicators for treatment, semester, and the covariates used for stratification during the randomization: academic department and instructor type (adjunct or full-time). In the remaining columns, we add controls for the following baseline characteristics reported in Table 1: student age, gender, race and credits accumulated; instructor contract and time preferences (using indicator variables for above/below median preference in the sample); and whether the course is a co-requisite.[26] All columns report standardized effect sizes except column 3, which reports effects in raw scores on a 0 to 100 point scale.

In all specifications, Instructor Incentives have an economically meaningful and statistically significant impact on student outcomes. In the full sample, Instructor Incentives improve student performance on the final exam by 0.20 SD ($p \approx 0.000$) or just over 6 percentage points off a control group mean of 52%. The estimated effects are smaller in the fall semester, 0.11 SD ($p = 0.048$), than in the spring semester, 0.25 SD ($p = 0.004$), though we cannot reject that these effect sizes are equal ($p = 0.61$). The pattern of effects also holds if we restrict the spring sample to the subset of instructors who also received incentives in the fall (Appendix Table A.3 column 7). These results

[25] Appendix Table A.3 column 1 shows that our results are robust to Ordinary Least Squares (OLS) estimation (i.e., omitting instructor-specific random effects).

[26] Our analysis differs from our pre-analysis plan in two ways. First, our pre-analysis plan includes GPA as a student-level covariate. We exclude GPA from our main analysis because it is missing for a substantial fraction of students. Second, our pre-analysis plan did not include co-requisite classification as a course-level covariate because we were not aware of this classification at the time. Columns 5 and 6 of Appendix Table A.3 repeat the analysis including GPA in two ways. In column 5, we impute GPA as the mean for students missing GPA and include an indicator for whether GPA is missing. In column 6, we run the analysis including GPA as a covariate and excluding students who are missing GPA. Neither specification affects the results. Column 4 of Appendix Table A.3 repeats our analysis excluding the covariate for co-requisite courses. No results are affected.

suggest that the effects of the incentives sustain themselves the first and second time they are offered.

We find no evidence that student incentives increase the effect of instructor incentives. The estimated effect of Combined Incentives is 0.15 SD ($p = 0.045$), which is smaller than the estimated effect of Instructor Incentives, though the two treatments are not statistically distinguishable. The estimated impact of Student Incentives alone is modest, 0.07 SD, and not statistically significant (as discussed Section 2.4 the Student Incentives treatment was insufficiently powered to detect meaningful effects). If the effects of Instructor Incentives and Student Incentives were additive, we would expect Combined Incentives to have a coefficient of 0.27 SD. The estimated impact of Combined Incentives of 0.15 SD therefore implies a subadditivity of -0.12 SD ($p = 0.231$), which is economically meaningful but not statistically significant (Appendix Table A.4). To address the concern that the estimated impact of Instructor Incentives is being buoyed by excluding less successful Combined Incentives sections, we also estimate effects on exam scores pooling Instructor Incentives with Combined Incentives. As shown in Appendix Table A.4, the estimated effects of pooled instructor incentives are large and statistically significant, 0.185 - 0.192 SD ($p \approx 0.000$).[27]

To explore the mechanisms behind the impact on exam scores, we separate the extensive margin, whether the student sat for the exam; and the intensive margin, the student's score conditional on taking the exam. These margins are presented in Table 3 for the full year sample. The extensive margin results report marginal effects from a random-effects probit regression. All specifications include the same student, instructor, and course covariates as in column 2 of Table 2.

Instructor Incentives have large and significant effects on both margins. At the extensive margin, Instructor Incentives increase rates of taking the exam by 5.3 percentage points ($p = 0.003$). Instructor Incentives also improve scores at the intensive margin, with exam scores improving an estimated 0.08 SD ($p = 0.040$) among the students who take the exam. This is particularly noteworthy given that the extensive margin result suggests that Instructor Incentives may induce more marginal students to take the exam, which could depress conditional exam scores. Turning to Combined Incentives, we find a large impact on the intensive margin of exam performance. Scores among those who take the exam increase by an estimated 0.13 SD ($p = 0.003$). How-

[27] Appendix Table A.4 presents the estimates pooling Instructor Incentives and Combined Incentives (column 1 pools Student Incentives with Control, column 2 includes a separate indicator for Student Incentives). Column 3 presents the results of the interaction specification used to estimate the subadditivity of Combined Incentives reported above. In column 4, limiting the interaction specification to the spring semester yields an estimated subadditivity of -0.15 SD ($p = 0.180$).

ever, there is no effect on the extensive margin of taking the exam. Student Incentives have little impact at either margin.

3.2 Course performance

One concern with using exam scores to measure improvement is that instructors may "teach to the test" in ways that do not improve (or even detract from) non-incentivized elements of the course. To address this, we explore the impact of incentives on overall course outcomes. Course grades were collected at the administrative level and were not directly incentivized. Thus, they provide a robustness check for our exam score results. Course grades depend partly on final exam scores and so are not independent. We discuss below the extent to which the impact on course grades would be expected if the final exam score was the only channel through which incentives affected grades.

Table 4 reports the impact of our treatments in the full year sample. We first estimate the impact of incentives on course grades in standardized units (column 1) and then grade points (column 2). We use the standard 0-4 scale of grade points corresponding to A-F grades with withdrawals counting as 0 grade points. We normalize grades within each department to have a mean of zero and a standard deviation of one. Column 3 reports the marginal probability that a student completes the course estimated using a random-effects probit regression. Column 4 presents the impact of incentives on the course grade conditional on course completion. All estimations use the same controls as column 2 of Table 2.

The effects of Instructor Incentives carry over to course outcomes, increasing course grades by 0.11 SD ($p = 0.039$) or 0.16 grade points ($p = 0.039$) off a control group mean of 2.08. Course completion rates increase by 3.7 percentage points ($p = 0.050$). This represents a 17% reduction in the baseline dropout rate of about 22%. As discussed above, course completion is a critical outcome for students, who receive no refund for the course if they withdraw and must meet a minimum completion rate to retain their financial aid. There is also a small increase at the intensive margin – course grades conditional on completion – that is not statistically significant. Combined Incentives have no impact on course completion but have positive, though not statistically significant, effects on course grades both unconditionally and conditional on completing the course, 0.07 and 0.09 SD, respectively. As in the exam results, we find little impact of Student Incentives (there is suggestive evidence of small negative effect on course grades).

The incentivized exams are worth between 5-25% of the course grade, depending on the course. Thus, if the exam is the only channel through which course grades are affected, then the 6 percentage point improvement on the exams found in Table 2 should translate to an improvement of $0.03 - 0.15$ grade points, depending on the course. Our estimated impact of 0.17 grade points lies above the upper bound of this interval, suggesting that our incentives are doing more than mechanically improving course performance though exam scores. In addition, the significant impact on course completion provides further evidence that the impact of incentives is not due to "gaming" on the exam itself.

3.3 Heterogeneity

We blocked our randomization on instructor classification: full-time or adjunct faculty. This guarantees balance across treatments for this characteristic and allows us to test for differential effects by instructor type – 48% of our students (48% of sections) are instructed by full-time faculty while 52% are instructed by adjunct faculty. Table 5 estimates the effects of incentives on exam scores, course grade and course completion by instructor type. Our specification includes indicator variables for instructor type, the full set of interactions of each instructor type with each treatment, and the full set of covariates. We also report p-values from tests of equality of the marginal impact of Instructor Incentives for full-time vs. adjunct faculty.

Our results suggest that there are heterogeneous effects by instructor type. Under Instructors Incentives, the exam scores and course grades of students taught by adjunct faculty improve by approximately 0.26 SD ($p \approx 0.000$) and 0.19 SD ($p = 0.003$), respectively. Instructor Incentives also increase course completion rates among adjunct faculty by an estimated 7 percentage points ($p = 0.007$). For full-time faculty, the effect of Instructor Incentives on exam scores is an estimated 0.13 SD, but not statistically significant ($p = 0.120$). There is no discernible impact on course grades or course completion. The larger effect of incentives on adjunct instructors compared to full-time faculty more than make up for the baseline difference in student performance we observe in the control group. The estimated effects of Combined Incentives and Student Incentives are similar to those in the full sample and do not appear to vary across instructor type.

We also blocked the randomization on department. Figure 1 presents the estimated within-department effects of Instructor Incentives on the normalized exam scores with 95% confidence intervals. The included departments are: Accounting (ACCT), Anatomy and Physiology (APHY), Art History (ARTH), Biology (BIOL), Business Operations Applications and Technology (BOAT), Business (BUSN), Communications (COMM), Criminology (CRIM), English (ENGL), Health Sciences (HLHS), Math

(MATH), Nursing (NRSG), Psychology (PSYC), and Software Development (SDEV).[28] While the small sample sizes within each department increase the error in our estimates, we find positive effects across the vast majority of departments. Psychology is the only department where results are even suggestively negative.

3.4 Dosage, spillovers, and substitution

While our primary effects are at the course level, we also explore the extent to which the impact of Instructor Incentives could scale up. A potential concern with our course-level results is that our treatment effects may be driven in part by instructors or students substituting effort away from unincentivized courses to give additional effort in the incentivized courses. If substitution is an important mechanism for our treatment effects, then incentives may have little or no impact when scaled up. In this section, we first examine substitution and returns to scale within courses in our program and then examine spillovers to courses outside our program. Our results suggest that instructor incentives could scale up without losing potency.

Table 6 presents the results within program courses. In order to focus our analysis on the impact of incentivizing instructors, we pool Instructor Incentives and Combined Incentives into "treatment" (the "control" group pools Student Incentives with Control).[29] We also report the p-value from a test of equality of treatment effects estimated within each panel. In Panel A, we estimate how the treatment effect changes as an instructor has more courses assigned to treatment.[30] In Panel B we similarly explore the returns to scale among students enrolled in multiple incentivized courses.[31] If substitution is driving our results, we would expect to find diminishing returns to scale. Treatment effects should be largest for instructors or students who can concentrate

[28] The Physical Science and Sociology departments did not have enough variation in treatment to estimate effects within department. The estimates include student-level covariates (age, gender, race, baseline credits). There was not enough variation within department to include semester fixed effects, instructor-level, or course-level covariates.

[29] We measure treatment at the semester level and report effects for the full sample. Regressions include semester fixed effects and all baseline covariates. As discussed above, the pooled treatment effect is 0.185 SD (Appendix Table A.4, column 1).

[30] We use an interacted random effects regression that controls for both the number of courses an instructor taught in the study and the number of credit hours the instructor taught outside of the study. Since all of an instructor's courses in the study are assigned to treatment or to control, we interact treatment assignment with each instructor type: A) one course in the study or B) multiple courses in the study.

[31] We use an interacted random effects regression that controls for the student's total number of courses both in and outside of the study. We interact treatment with each student dosage type: A) one course exposed to incentives or B) multiple courses exposed to incentives. Students exposed only to control courses are the omitted group.

their effort on only one incentivized course compared to instructors or students who have to spread their effort across multiple incentivized courses. We find no evidence that the impact of incentives decreases with either instructor or student dosage. In fact, for both instructors and students the point estimates are larger (though statistically indistinguishable) for those exposed to instructor incentives in multiple courses than for those treated in only one course.

In Panel C, we look for evidence that students may be substituting effort between incentivized and unincentivized courses. We estimate treatment effects for students who have courses both in treatment and in control, and for students whose courses are only in treatment or only in control.[32] The first set of students has scope to substitute effort away from unincentivized courses towards incentivized courses. If this occurs, it should increase treatment effects – by lowering the control course outcomes, increasing the treatment course outcomes, or both – compared to the second set of students who has no opportunity to substitute. As in the first two panels, we find no evidence for substitution. The estimated impact of incentives is larger for students whose courses are all in the same treatment and thus have no scope for substitution (the difference in the point estimates is large but not statistically significant).

We next examine spillovers of incentives to outcomes outside our program in Table 7. Columns 1 and 2 report estimates of the impact of incentives on grades (measured in grade points) and course completion in all concurrent, non-program courses. The instructors for these courses either did not choose to participate or the courses were ineligible because they lacked an objective course exam. In columns 3 and 4, we estimate the effects of each treatment on students' cumulative GPA and credit accumulation. These measures include all courses taken prior to and during the treatment semester. In column 5, we estimate the impact on earning a degree, transferring to a four-year college or enrolling in the subsequent semester post-treatment. These estimates are at the student-course level and thus capture the average impact on overall outcomes of incentivizing a single course.[33]

The estimated effects of Instructor Incentives are generally small and positive. We

[32]We use an interacted random effects regression with controls for the student's total number of courses in- and outside of the study. We include indicator variables for each student type: A) all courses in one treatment or B) courses split between incentives and control. We report the interaction of each student type with treatment.

[33]Because the experiment was randomized at the instructor-section level rather than the student level, these estimates also "double count" the outcomes of students exposed to incentives (or control) in multiple courses. All estimates pool the full year sample and use random effects estimation with the full set of baseline covariates (columns 2 and 5 use random effects probit estimation). The dependent variable is an indicator equal to one if a student graduates (8.4% of observations), transfers to a four year school (2.8% of observations) or enrolls (71.33% of observations).

find evidence of positive spillovers to course completion in non-program courses of 3.0 percentage points ($p = 0.010$). These effects are only slightly smaller than the 3.9 percentage point increase in completion in the incentivized course (Table 4). This finding suggests that Instructor Incentives may be leading students to stay enrolled in school more generally during the treatment semester. The small impact on cumulative outcomes and post-treatment enrollment is perhaps not surprising given that we are examining the impact of incentivizing a single course and the median student is enrolled in four courses each semester. One desirable test of treatment persistence is the impact of an instructor in an introductory course on subsequent, follow-on courses in an academic track. Unfortunately, the diverse set of courses in our study includes courses that are not introductory and courses that are not part of set tracks, which makes this test less applicable in our sample.[34]

3.5 Expenditures and time use

A potential mechanism for the impact of Instructor Incentives on student performance is through effects on instructors' financial or time expenditures. For example, providing instructors with money at the beginning of the semester could allow them to purchase resources that aid in teaching. Or, the additional money could reduce the need for outside employment increasing the amount of time instructors can devote to preparing class and supporting students outside of class. This is particularly relevant for community college instructors who have high rates of outside employment (83% of our sample report some amount of outside employment).

Table 8 presents means from self-reported instructor surveys about personal money spent on course materials or professional development (Panel A); and time spent during a typical week in the semester on both teaching-related activities and outside employment (Panel B).[35] The differences between treatment and control instructors are generally small and not statistically significant. The lack of meaningful differences in time use and expenditures suggests that the impact of incentives on performance may

[34] We separately examine treatment effects on summer enrollment and find that offering the summer scholarship incentive to students increases summer enrollment by 4.6 percentage points ($p = 0.080$) in Student Incentives and 7.1 percentage points ($p = 0.037$) in Combined Incentives. The Instructor Incentives treatment, which did not offer summer scholarships, does not have a significant impact on summer enrollment. Results available upon request

[35] The survey outcome is reported for each row. Instructors report expenditures for a \$0-\$500 range and time use for a range of 0-16 hours in each category. Half of the responses for outside employment are top coded at 16 hours. Columns 1 and 2 present the means for control and treatment instructors, respectively. We report tests of statistical significance from a random-effects regression with controls for semester and instructor type and standard errors clustered at the instructor level.

result from more subtle changes to teaching that are not captured by hours spent on particular tasks (or could be due to the difficulty of accurately measuring time use).

We also collected administrative data for course evaluations completed by students in every course and find no evidence of treatment effects on a series of measures such as "I would recommend this instructor to others" (Appendix Table A.5). We note that response rates were low (fewer than a third of students completed the evaluations) and the measures are heavily top-coded with a median rating of 5 out of 5.

4 Contract Preferences

As a personnel policy, loss contracts need to be not only effective but also palatable to employees (instructors, in our case). Accordingly, we examine instructors' preferences for the loss contracts we offer. As discussed above, standard behavioral models predict that people will prefer to work under gain contracts rather than loss contracts. In practice, there is limited empirical evidence on employee preferences between such contracts.

4.1 Baseline preferences

When instructors enrolled in the study (either before the Fall semester or before the Spring semester), we used the incentive-compatible multiple price list mechanism described in Section 2.3 to elicit their baseline preferences between loss and gain contracts. For instructors who participated in the Fall semester, we also elicited their contract preferences at the end of the Fall semester (we could not incentivize endline preferences for Spring semester instructors because we did not provide incentives after the Spring semester).

From the multiple price list, instructors revealed the price ratio at which they preferred to receive the loss contract, which provides upfront bonuses, rather than a gain contract, which awards bonuses at the end of the semester. We then estimate the per-student bonus amount that an instructor is willing to sacrifice in order to receive a loss contract, where positive values indicate a preference for loss contracts and negative values indicate a preference for gain contracts – i.e., an instructor needs to be paid a higher per-student bonus to work under a loss contract rather than gain contract.

Figure 2 plots the histogram of baseline contract preferences elicited on the initial enrollment survey (either Fall or spring). Preferences are calculated using midpoint estimation where possible.[36] We find a preference for gain contracts at baseline: on average, instructors prefer gain contracts until loss contracts offer \$4.58 more per student, which is equivalent to 9.16% of the \$50 per student incentives. For the average instructor, this represents a *potential* difference in incentive payments between the loss and gain contracts of \$138 and an *expected* difference of \$66 using average pass rates in the treatment group.[37] We estimate that this willingness to pay to avoid loss contracts corresponds to a loss aversion parameter of 1.99, which is in line with the literature from laboratory experiments (e.g. Tversky and Kahneman, 1991; Abdellaoui et al., 2008).[38] Using interval estimation rather than midpoint estimation, average WTP for a gain contract is \$6.13 per student (12.3% of the \$50 per student incentive), yielding $\lambda = 2.63$.

Approximately half of respondents reveal the strongest preference for gain contracts (i.e., the minimum value), preferring gain contracts even if they offer \$10 less per student than the loss contract. For the average instructor, this represents a potential difference of \$301 and an expected difference of \$144, corresponding to an estimated loss aversion parameter of at least 4.35. Such choices thus indicate substantial loss aversion. Or, alternatively, could reflect confusion about the loss contract or lack of attention when taking the survey (i.e., filling in the same contract choice for every decision).

The second most common response, 31% of respondents, is a weak preference for gain contracts. These instructors prefer gain contracts if per-student bonuses are equal but will switch to preferring a loss contract if the gain contract offers \$1 less per student. We categorize these instructors' contract values as $-\$.50$, which is the midpoint between $-\$1$ (when the instructor prefers the loss contract) and \$0 (when the

[36]When instructors switched only once across the multiple price list, we assigned them a value equal to the midpoint of the interval over which their preferences shifted. When they never switched, we assigned them the minimum or maximum value from the list, ensuring that their assigned preferences exceeded anyone who switched in the interior. We dropped anyone who switched multiple times in the list. This drops 1 instructor in the Fall baseline and 3 in the Fall endline.

[37]At a difference of \$4.58 per student, the lower expected incentive payment under the gain contract is $\$4.58/student \times \Pr[Pass] \times N = \$4.58 \times 0.478 \times 30.1 = \$66$ where $\Pr[Pass] = 0.478$ is the pass rate in the treatment group and $N = 30.1$ is the average number of students per instructor per semester.

[38]We calculate the loss aversion parameter, λ, such that the lower expected incentive payment under the gain contract equals the disutility from losses under the loss contract. We assume that instructors have rational expectations and anticipate a pass rate equal to the observed mean rate among treatment instructors; loss and gain contracts have identical motivating effects; there is no discounting between the payment dates; and disutility from losses equals $\lambda(x - r)$ for $x < r$ where $x = \Pr[Pass](\$50N)$ is the expected incentive payment for an instructor with N students and $r = 0.5(\$50N)$ is the reference point (i.e., the upfront bonus). Thus, for indifference between a \$50/student loss contract and a \$45.4/student gain contract, an average of $N = 30.1$ students, and an expected pass rate $\Pr[Pass] = 0.478$, we set $\lambda(.478 \times \$50 \times 30.1 - .5 \times \$50 \times 30.1) = -\$66$.

instructor prefers the gain contract). This corresponds to a loss aversion parameter between 0.43 and 0 (i.e., no reference dependence in preferences).

4.2 Effects of experience with incentives on preferences

We next examine the effect of treatment on instructor preferences. Figure 3 summarizes the changes in instructor preferences by treatment group between Fall baseline and Fall endline measured in the enrollment and mid-year surveys, respectively.[39] The modal instructor in both the treatment and control groups has no change in preferences. However, treatment instructors are more likely to change their preference and do so towards preferring loss contracts: 38.0% of treatment instructors show increased preference towards loss contracts compared to only 12.8% of control instructors. Of the treatment instructors who change their preferences, 28.6% move from the strongest preference for gain contracts to a weak preference for gain contracts, making it the most common shift in preferences.[40]

In Table 9, we estimate the effect of experience with loss contracts on contract preferences. In the analysis below, we use an interval regression to correct for the interval censoring from our multiple price list. The outcome variable in column 1 is preference for loss contracts in the Fall baseline survey – i.e., a test of the balance of baseline preferences between treatment and control instructors (we restrict our sample to the instructors for whom we have both baseline and endline preferences). The outcome variable in columns 2-4 is contract preference in the Fall endline survey. Columns 1-2 only control for instructor type. Column 3 adds controls for baseline contract preferences.[41] Column 4 additionally controls for the instructor's baseline discount rate. We report the p-value from a test of whether the treatment group's value for the loss contract is equal to \$0 – that is, if treatment instructors are indifferent between loss and gain contracts of equal value.

Column 1 demonstrates that there are no baseline differences in contract preferences between treatment and control instructors. Both groups significantly prefer gain contracts: the average instructor would need to receive a little over \$6 *more* per student in order to prefer the loss contract ($p < 0.01$ for both the treatment and control

[39] Unlike in the analysis of baseline preferences, we do not include spring instructors in the analysis of changes in preferences because, as noted above, the end of the study meant that contract choices on the spring endline survey could not be incentive-compatible.

[40] The individual-level changes in preferences for treatment and control instructors are plotted in Figures A.3a and A.3b.

[41] We use dummy variables for five categories of switching points: less than -\$10 (i.e., always prefers the gain contract), -\$10 to -\$0.5, -\$0.5 to \$0, \$0 to \$10, and greater than \$10 (always prefers the loss contract). Group 3: -\$0.5 to \$0 (i.e., the median instructor) is the omitted group.

group).

Columns 2-4 estimate the impact of receiving (loss-framed) Instructor Incentives during the fall semester. As at baseline, Control instructors continue to prefer gain contracts by a little over \$6 ($p = 0.001$). In contrast, assignment to Instructor Incentives significantly increases instructor preferences for loss contracts. The treatment effects of \$4.33 - \$4.53 ($p < 0.05$ in all specifications) largely erase preferences for gain contracts. In the endline survey, the treatment group's value for the loss contract is no longer statistically distinguishable from zero ($p > 0.20$ in all specifications). That is, after experiencing loss incentives, instructors become (close to) indifferent between the two contract types.

4.3 Well-being and stress

Finally, we examine the impact of Instructor Incentives on well-being and stress. On the one hand, upfront cash bonuses may relieve financial burdens and stress. On the other hand, the anxiety of potentially having to pay back the bonus could induce stress. If the latter concern dominates, then loss contracts may lower employee welfare and increase employee turnover. To address this concern, we surveyed instructors about their personal and professional well-being. Table 10 reports survey response means for the Control and Instructor Incentive groups (measured on a 1-10 scale where 10 is the highest agreement with the survey statement). The differences between treatment and control instructors are generally small and not statistically significant.[42] We find little evidence that the incentives meaningfully affect well-being or stress, or that instructors who work under them indicate a higher likelihood of leaving their job.

5 Conclusion

Ours is the first study to test the impact of an intervention aimed at college instruction on student performance. We demonstrate that performance-based incentives for community college instructors have a large impact on student outcomes, equivalent to improving instructor quality by a standard deviation. At an expected cost of \$25 per student-course, instructor incentives represent a relatively low-cost option for improving student performance and encouraging student retention, both critical outcomes for community colleges.

[42]We conducted tests of significant differences between groups using random effects regressions including controls for semester, instructor type and baseline response to the same question in the enrollment survey.

Interestingly, we find no evidence that combining instructor incentives with student incentives increases their effectiveness. If anything, we find evidence that combined incentives are *less* effective. While only suggestive, our findings could perhaps be due to strategic substitution. If instructors believe that their own effort and the effort of students are substitutes, they may have reduced their effort in classes that received student incentives (e.g., encouragement to continue the course) believing that increased student effort would compensate. More broadly, our results point to the need for future research on the interaction between student and instructor incentives.

Our study also provides support for the use of loss-framed contracts as a policy tool. The incentives are feasible and effective across a broad range of departments. Their impact does not diminish over semesters, suggesting that the motivational power of loss-framed incentives can be sustained.[43] Similarly, the effects are constant (if not increasing) as instructors and students are exposed to more courses. We also find evidence that incentives encourage general enrollment during the treatment semester, increasing completion rates not only in incentivized courses but also in courses outside our program. Finally, experience with the contracts increases instructors' desire to work under them. This novel finding may be due to instructors learning that working under loss-framed incentives is less painful (or more beneficial) than they expected, or could result from increased familiarity with these unusual contracts.

Taken together, our findings suggest that loss-framed incentives could be scaled up and incorporated into personnel contracts. Scaling up instructor incentives across courses would allow for an examination of their impact on student persistence and educational attainment. Such scale up could also examine institutional costs over and above incentive payments, which were not relevant in our study, such as developing standardized exams and incorporating incentives into payroll systems.

Instructor incentives may be particularly relevant for two year colleges and other non-selective institutions. The primary focus of their instructors is teaching (rather than research). They have lower average salaries than four year selective colleges, so that equivalent incentives represent a larger percentage relative to instructors' base pay (Miller and Topper, 2017). And, as discussed above, the prior work on instructor

[43] The limited prior work on the impact of loss-framing over time is mixed. List (2003, 2004, 2011) finds that experience limits the impact of loss framing in trading markets. In contrast, Hossain and List (2012) conduct an experiment offering incentives to Chinese factory workers and find that the effects of loss-framing are sustained over time. Closest to our study, Fryer Jr et al. (2018) find that while upfront bonuses for teachers have large impacts in the first year they are offered, the incentives are not effective in the second year of their experiment. We note that they re-randomize in the second year of their experiment so that teachers receive different treatments across years and there is no group of never incentivized teachers across both years.

quality suggests that improving instructor effectiveness at non-selective institutions may have a particularly large impact on student performance.

Performance-based incentives could also be targeted towards adjunct instructors at both two year and four year colleges. Adjuncts receive lower salaries than full-time faculty, which may drive the larger treatment affects among adjuncts in our study. Adjuncts also work under short term flexible contracts that could potentially be restructured. More broadly, personnel policies for adjunct instructors are of critical importance. They comprise a growing share of instructors at postsecondary institutions, representing about half of all instructors and nearly eighty percent of instructors at public community colleges (Hurlburt and McGarrah, 2016).

The purpose of this study was to demonstrate that incentives can improve instructor effectiveness at the postsecondary level. We limited our focus to loss-framed incentives, but it may be the case that gain-framed incentives could also be effective in this context. Because of instructors' baseline preferences as well as logistical concerns – for example, collecting repayments – gain-framed contracts are potentially preferable. However, it could also be the case that loss-framed contracts serve as a commitment device that instructors learn to prefer because they anticipate working harder under them. Given their demonstrated effectiveness and low cost, we believe that future work is warranted on the optimal design and implementation of incentive contracts for college instructors.

References

Abdellaoui, M., Bleichrodt, H., and L'Haridon, O. (2008). A tractable method to measure utility and loss aversion under prospect theory. *Journal of Risk and uncertainty*, 36(3):245.

Armantier, O. and Boly, A. (2015). Framing of incentives and effort provision. *International Economic Review*, 56(3):917–938.

Bailey, T. R., Jaggars, S. S., and Jenkins, D. (2015). *Redesigning America's community colleges*. Harvard University Press.

Barlevy, G. and Neal, D. (2012). Pay for percentile. *The American Economic Review*, 102(5):1805–1831.

Barrera-Osorio, F. and Raju, D. (2017). Teacher performance pay: Experimental evidence from pakistan. *Journal of Public Economics*, 148:75–91.

Barrow, L., Richburg-Hayes, L., Rouse, C. E., and Brock, T. (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.

Behrman, J. R., Parker, S. W., Todd, P. E., and Wolpin, K. I. (2015). Aligning learning incentives of students and teachers: Results from a social experiment in mexican high schools. *Journal of Political Economy*, 123(2):325–364.

Bettinger, E., Fox, L., Loeb, S., and Taylor, E. (2014). Changing distributions: How online college classes alter student and professor performance. Technical report, Working Paper, Stanford University.

Bettinger, E. P. and Long, B. T. (2006). The increasing use of adjunct instructors at public institutions: Are we hurting students. *Whats happening to public higher education*, pages 51–69.

Bettinger, E. P. and Long, B. T. (2010). Does cheaper mean better? the impact of using adjunct instructors on student outcomes. *The Review of Economics and Statistics*, 92(3):598–613.

Braga, M., Paccagnella, M., and Pellizzari, M. (2016). The impact of college teaching on students academic and labor market outcomes. *Journal of Labor Economics*, 34(3):781–822.

Brodaty, T. and Gurgand, M. (2016). Good peers or good teachers? evidence from a french university. *Economics of Education Review*, 54:62–78.

Brooks, R. R., Stremitzer, A., and Tontrup, S. (2012). Framing contracts: Why loss framing increases effort. *Journal of Institutional and Theoretical Economics JITE*, 168(1):62–82.

Carrell, S. E. and West, J. E. (2010). Does professor quality matter? evidence from random assignment of students to professors. *Journal of Political Economy*, 118(3):409–432.

Courant, P. N. and Turner, S. (2017). Faculty deployment in research universities. Technical report, National Bureau of Economic Research.

De Fraja, G., Oliveira, T., and Zanchi, L. (2010). Must try harder: Evaluating the role of effort in educational attainment. *The Review of Economics and Statistics*, 92(3):577–597.

De Quidt, J. (2017). Your loss is my gain: a recruitment experiment with framed incentives. *Journal of the European Economic Association*.

De Quidt, J., Fallucchi, F., Kölle, F., Nosenzo, D., and Quercia, S. (2017). Bonus versus penalty: How robust are the effects of contract framing? *Journal of the Economic Science Association*, 3(2):174–182.

De Vlieger, P., Jacob, B., and Stange, K. (2017). Measuring instructor effectiveness in higher education. In *Productivity in Higher Education*. University of Chicago Press.

DellaVigna, S. and Pope, D. (2016). What motivates effort? evidence and expert forecasts. *The Review of Economic Studies*.

Deming, D. and Dynarski, S. (2009). Into college, out of poverty? policies to increase the postsecondary attainment of the poor. Technical report, National Bureau of Economic Research.

Duflo, E., Hanna, R., and Rya, S. P. (2012). Incentives work: Getting teachers to come to school. *The American Economic Review*, 102(4):1241–1278.

Ehrenberg, R. G. (2012). American higher education in transition. *The Journal of Economic Perspectives*, 26(1):193–216.

Ehrenberg, R. G. and Zhang, L. (2005). Do tenured and tenure-track faculty matter? *Journal of Human Resources*, 40(3):647–659.

Englmaier, F., Grimm, S., Schindler, D., and Schudy, S. (2018). Effect of incentives in non-routine analytical team tasks – evidence from a field experiment.

Evans, W. N., Kearney, M. S., Perry, B. C., and Sullivan, J. X. (2017). Increasing community college completion rates among low-income students: Evidence from a randomized controlled trial evaluation of a case management intervention. Technical report, National Bureau of Economic Research.

Figlio, D. N., Schapiro, M. O., and Soter, K. B. (2015). Are tenure track professors better teachers? *Review of Economics and Statistics*, 97(4):715–724.

Fryer, R. G. (2013). Teacher incentives and student achievement: Evidence from new york city public schools. *Journal of Labor Economics*, 31(2):373–407.

Fryer, R. G. (2017). The production of human capital in developed countries: Evidence from 196 randomized field experimentsa. *Handbook of Economic Field Experiments*, 2:95–322.

Fryer Jr, R. G., Levitt, S. D., List, J., and Sadoff, S. (2012). Enhancing the efficacy of teacher incentives through loss aversion: A field experiment. Technical report, National Bureau of Economic Research.

Fryer Jr, R. G., Levitt, S. D., List, J., and Sadoff, S. (2018). Enhancing the efficacy of teacher incentives through loss aversion: A field experiment. Technical report, Working paper.

Geng, T. (2017). The complementarity of incentive policies in education: Evidence from new york city.

Glazerman, S., McKie, A., and Carey, N. (2009). An evaluation of the teacher advancement program (tap) in chicago: Year one impact report. final report. *Mathematica Policy Research, Inc.*

Glewwe, P., Ilias, N., and Kremer, M. (2010). Teacher incentives. *American Economic Journal: Applied Economics*, 2(3):205–227.

Hoffmann, F. and Oreopoulos, P. (2009). Professor qualities and student achievement. *The Review of Economics and Statistics*, 91(1):83–92.

Hong, F., Hossain, T., and List, J. A. (2015). Framing manipulations in contests: a natural field experiment. *Journal of Economic Behavior & Organization*, 118:372–382.

Hossain, T. and List, J. A. (2012). The behavioralist visits the factory: Increasing productivity using simple framing manipulations. *Management Science*, 58(12):2151–2167.

Hurlburt, S. and McGarrah, M. (2016). The shifting academic workforce: Where are the contingent faculty. *Delta Cost Project at American Institutes for Research. Retrieved from http://www. air. org/sites/default/files/downloads/report/Shifting-Academic-Workforce-November-2016. pdf.*

Imas, A., Sadoff, S., and Samek, A. (2016). Do people anticipate loss aversion? *Management Science*, 63(5):1271–1284.

Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the econometric society*, pages 263–291.

Kermer, D. A., Driver-Linn, E., Wilson, T. D., and Gilbert, D. T. (2006). Loss aversion is an affective forecasting error. *Psychological science*, 17(8):649–653.

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

Lavecchia, A. M., Liu, H., and Oreopoulos, P. (2014). Behavioral economics of education: Progress and possibilities. Technical report, National Bureau of Economic Research.

Levitt, S. D., List, J. A., Neckermann, S., and Sadoff, S. (2016). The behavioralist goes to school: Leveraging behavioral economics to improve educational performance. *American Economic Journal: Economic Policy*, 8(4):183–219.

List, J. A. (2003). Does market experience eliminate market anomalies? *The Quarterly Journal of Economics*, 118(1):41–71.

List, J. A. (2004). Neoclassical theory versus prospect theory: Evidence from the marketplace. *Econometrica*, 72(2):615–625.

List, J. A. (2011). Does market experience eliminate market anomalies? the case of exogenous market experience. *The American Economic Review*, 101(3):313–317.

List, J. A., Livingston, J. A., and Neckermann, S. (2012). Harnessing complementarities in the education production function. Technical report, Working Paper, available at https://econresearch. uchicago. edu/sites/econresearch. uchicago. edu/files/Livingston_AWFE% 20version. pdf.

List, J. A. and Samek, A. S. (2015). The behavioralist as nutritionist: leveraging behavioral economics to improve child food choice and consumption. *Journal of health economics*, 39:135–146.

Loyalka, P. K., Sylvia, S., Liu, C., Chu, J., and Shi, Y. (2016). Pay by design: Teacher performance pay design and the distribution of student achievement.

Mbiti, I., Muralidharan, K., Romero, M., Schipper, Y., Manda, C., and Rajani, R. (2016). Inputs, incentives, and complementarities in primary education: Experimental evidence from tanzania. *Unpublished Paper*.

McFarland, J., Hussar, B., de Brey, C., Snyder, T., Wang, X., Wilkinson-Flicker, S., Gebrekristos, S., Zhang, J., Rathbun, A., Barmer, A., et al. (2017). The condition of education 2017. nces 2017-144. *National Center for Education Statistics*.

Miller, A. and Topper, M. (2017). Trends in faculty salaries: 199596 to 201415. Technical report, National Education Association.

Muralidharan, K. and Sundararaman, V. (2011). Teacher performance pay: Experimental evidence from india. *Journal of political Economy*, 119(1):39–77.

NCCBP (2014). *National Community College Benchmark Project Report*.

Neal, D. (2011). The design of performance pay in education. Technical report, National Bureau of Economic Research.

Oreopoulos, P. and Petronijevic, U. (2013). Making college worth it: A review of the returns to higher education. *The Future of Children*, 23(1):41–65.

Patel, R., Richburg-Hayes, L., de la Campa, E., and Rudd, T. (2013). Performance-based scholarships: What have we learned? interim findings from the pbs demonstration.

Ran, X. and Xu, D. (2016). How and why do adjunct instructors affect students academic outcomes? evidence from two-year and four-year colleges.

Rogers, G. S. (2015). Part-time faculty and community college student success. *Community College Journal of Research and Practice*, 39(7):673–684.

Sadoff, S. (2014). The role of experimentation in education policy. *Oxford Review of Economic Policy*, 30(4):597–620.

Shapiro, D., D. A. H. F. W. P. Y. X. N. A. . H. Y. A. (2017). Completing college: A national view of student attainment rates (signature report no. 12b). *National Student Clearinghouse*.

Springer, M. G., Ballou, D., Hamilton, L., Le, V.-N., Lockwood, J., McCaffrey, D. F., Pepper, M., and Stecher, B. M. (2011). Teacher pay for performance: Experimental evidence from the project on incentives in teaching (point). *Society for Research on Educational Effectiveness*.

Springer, M. G., Pane, J. F., Le, V.-N., McCaffrey, D. F., Burns, S. F., Hamilton, L. S., and Stecher, B. (2012). Team pay for performance: Experimental evidence from the round rock pilot project on team incentives. *Educational Evaluation and Policy Analysis*, 34(4):367–390.

Tversky, A. and Kahneman, D. (1991). Loss aversion in riskless choice: A reference-dependent model. *The quarterly journal of economics*, 106(4):1039–1061.

Figure 1: Effects of Instructor Incentives on Normalized Exam Scores by Department

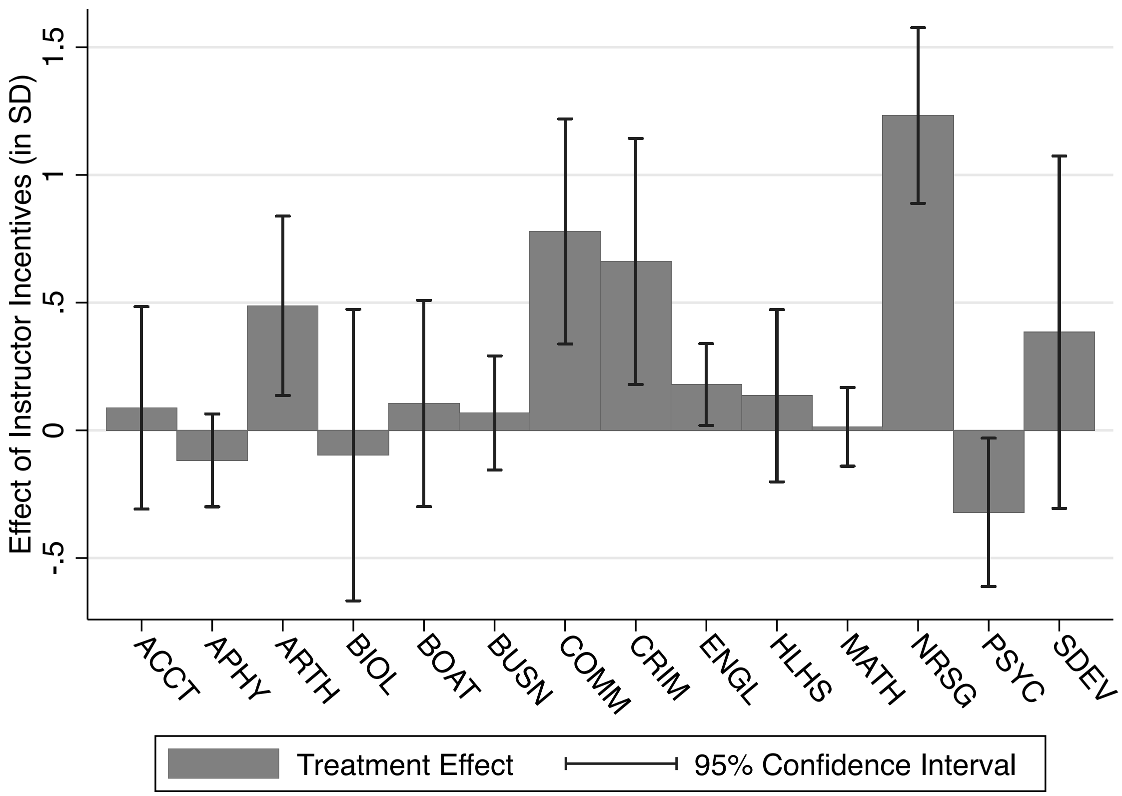

Notes: The figure presents coefficients and 95% confidence intervals for Instructor Incentives from random effects linear estimation with standard errors clustered by instructor within the following departments: Accounting (ACCT), Anatomy and Physiology (APHY), Art History (ARTH), Biology (BIOL), Business Operations Applications and Technology (BOAT), Business (BUSN), Communications (COMM), Criminology (CRIM), English (ENGL), Health Sciences (HLHS), Math (MATH), Nursing (NRSG), Psychology (PSYC), and Software Development (SDEV). The dependent variable is exam score standardized within department. All regressions include student-level covariates (age, gender, race/ethnicity, baseline credits).

Figure 2: Baseline Contract Preferences

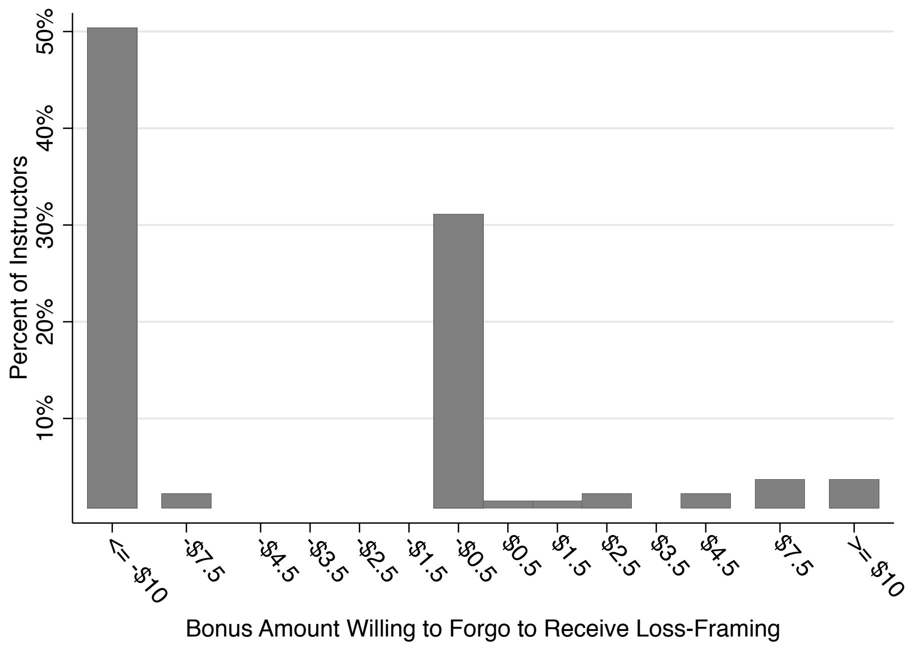

Notes: The figure presents the distribution of baseline contract preferences for all instructors. < -$10 indicates instructors who preferred a gain-framed bonus of $40 per student over a loss-framed bonus of $50 per student. > $10 indicates instructors who preferred a loss-framed bonus of $50 per student over a gain-framed bonus of $60 per student. All other values are the mid-point between the per-student bonus amounts over which the instructor "switches" from preferring the gain-framed bonus to preferring the loss-framed bonus.

Figure 3: Change in Contract Preferences by Treatment

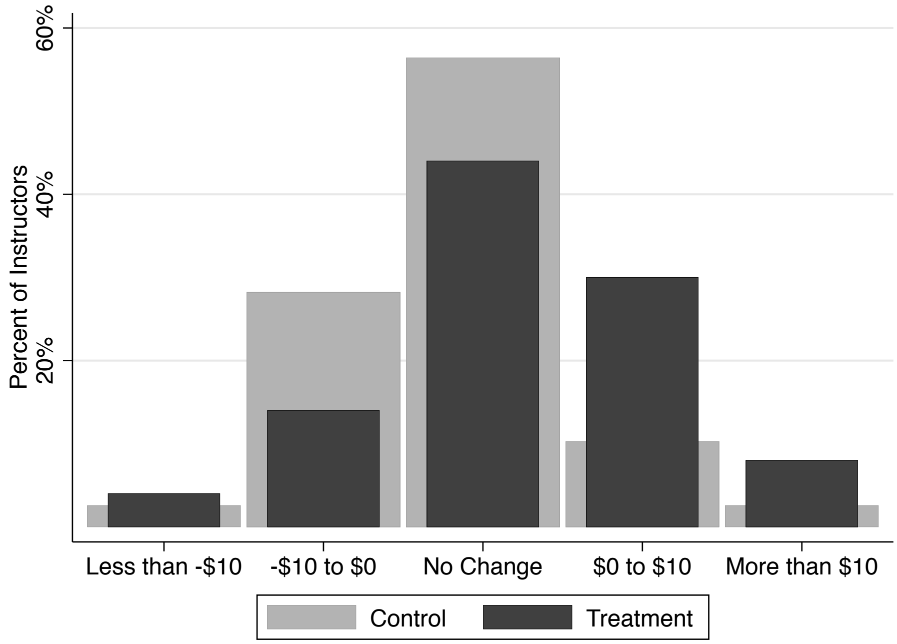

Notes: The figure presents the distribution of instructors' changes in willingness to pay for loss-framed contracts between the fall baseline and fall endline surveys. Positive values indicate increased willingness to pay for loss-framed contracts.

Table 1: Baseline Characteristics by Treatment and Semester

	Fall			Spring				
	Control	Instructor Incentives	**Fall F-test**	Control	Instructor Incentives	Combined Incentives	Student Incentives	**Spring F-test**
Student-Level Characteristics								
Age	24.385	24.594	**0.754**	25.228	24.758	24.787	24.659	**0.835**
	(0.423)	(0.516)		(0.473)	(0.545)	(0.613)	(0.642)	
Male	0.342	0.340	**0.944**	0.326	0.274	0.322	0.325	**0.474**
	(0.028)	(0.023)		(0.034)	(0.029)	(0.026)	(0.044)	
White	0.817	0.818	**0.964**	0.817	0.831	0.830	0.859*	**0.328**
	(0.011)	(0.015)		(0.015)	(0.015)	(0.018)	(0.017)	
Baseline Credits	12.403	13.018	**0.752**	18.011	19.442	19.236	17.580	**0.883**
	(1.012)	(1.665)		(1.442)	(2.525)	(1.315)	(1.986)	
Baseline GPA	2.843	2.804	**0.529**	2.934	2.925	2.950	2.790*	**0.254**
	(0.037)	(0.050)		(0.048)	(0.068)	(0.056)	(0.066)	
Missing GPA	0.372	0.371	**0.981**	0.498	0.446	0.416*	0.514	**0.273**
	(0.030)	(0.030)		(0.028)	(0.039)	(0.035)	(0.041)	
Instructor-Level Characteristics								
Full-Time	0.377	0.345	**0.733**	0.442	0.514	0.378	0.368	**0.284**
	(0.067)	(0.065)		(0.077)	(0.083)	(0.081)	(0.114)	
Total Sections	1.981	1.727	**0.273**	2.093	2.000	2.000	2.211	**0.898**
	(0.190)	(0.131)		(0.182)	(0.182)	(0.182)	(0.224)	
Students Per Section	16.467	16.342	**0.906**	14.200	15.893	15.164	14.421	**0.520**
	(0.749)	(0.750)		(0.709)	(0.866)	(0.697)	(1.042)	
Below-Median	0.509	0.455	**0.572**	0.558	0.378	0.459	0.684	**0.151**
Contract Value	(0.069)	(0.068)		(0.077)	(0.081)	(0.083)	(0.110)	
Below-Median	0.642	0.691	**0.591**	0.651	0.757	0.676	0.632	**0.523**
Discount Rate	(0.067)	(0.063)		(0.074)	(0.072)	(0.078)	(0.114)	
Section-Level Characteristics								
Co-Requisite	0.276	0.263	**0.905**	0.056	0.068	0.044	0.000	**0.542**
	(0.079)	(0.075)		(0.055)	(0.065)	(0.044)	(0.000)	
Observations	1776	1592		1043	657	687	311	

Table reports means/proportions for each group with standard errors in parentheses, are clustered at the level of randomization (instructor).
p-values are reported for joint orthogonality test across treatment groups.
Significance tests against control mean: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2: Effects of Incentives on Exam Scores

	Pooled Fall and Spring Semesters (Standardized Score)		(Raw Score)	Fall Semester	Spring Semester
Instructor Incentives	0.202***	0.204***	6.087***	0.113**	0.247***
	(0.057)	(0.056)	(1.794)	(0.057)	(0.085)
Combined Incentives	0.154**	0.148**	4.637*		0.166**
	(0.078)	(0.075)	(2.452)		(0.084)
Student Incentives	0.070	0.065	1.815		0.067
	(0.074)	(0.076)	(2.498)		(0.083)
Department	Yes	Yes	Yes	Yes	Yes
Instructor Type	Yes	Yes	Yes	Yes	Yes
Baseline Characteristics	No	Yes	Yes	Yes	Yes
Pr(Instructor = Combined)	0.506	0.424	0.527		0.268
Instructors	127	127	127	100	96
Observations	5839	5839	5839	3189	2650

Random effects linear estimation.

Standard errors in parentheses clustered at the instructor level.

Dependent variable: exam score standardized within dept. (mean 0, s.d. 1) except in column 3.

Column 1-3 include semester fixed effects.

Column 2-5 include covariates for student (age, gender, race/ethnicity, baseline credits), instructor (contract value, discount rate) and course (co-requisite).

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3: Effects of Incentives on Extensive and Intensive Margins

	Take Exam	Score if Take
Instructor Incentives	0.053***	0.082**
	(0.018)	(0.040)
Combined Incentives	0.010	0.135***
	(0.029)	(0.045)
Student Incentives	0.012	0.054
	(0.027)	(0.047)
Control Group Mean	0.746	0.451
	(0.44)	(0.51)
Pr(Instructor = Combined)	0.124	0.081
Instructors	127	127
Observations	5839	4419

Column 1: Marginal effects from random effects probit estimation.

Column 2: Random effects linear estimation.

Standard errors clustered at the instructor level in parentheses.

Standard deviation reported for control group mean.

Dependent variable column 2: exam score standardized within dept. (mean 0, s.d. 1).

Includes semester and department fixed effects and covariates for student (age, gender, race/ethnicity, baseline credits), instructor (type, contract value, discount rate) and course (co-requisite).

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Effects of Incentives on Course Performance

	Course Grade (Standardized)	Course Grade (Grade Points)	Complete Course	Grade if Complete
Instructor Incentives	0.108**	0.165**	0.037**	0.037
	(0.052)	(0.080)	(0.019)	(0.051)
Combined Incentives	0.071	0.113	-0.010	0.094
	(0.072)	(0.113)	(0.029)	(0.063)
Student Incentives	-0.058	-0.100	-0.002	-0.068
	(0.051)	(0.078)	(0.025)	(0.048)
Control group mean	-0.02	2.08	0.783	0.343
	(0.99)	(1.55)	(0.41)	(0.80)
Pr(Instructor = Combined)	0.560	0.603	0.071	0.165
Instructors	134	134	134	134
Observations	6066	6066	6066	4797

Column 1, 2 and 4: Random effects linear estimation.

Column 3: Marginal effects from random effects probit estimation.

Standard errors in parentheses clustered at the instructor level.

Standard deviation reported for control group mean.

Dependent variable column 2 and 4: course grade standardized within dept. (mean 0, s.d. 1).

Includes semester and department fixed effects and covariates for student (age, gender, race/ethnicity, baseline credits), instructor (type, contract value, discount rate) and course (co-requisite).

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5: Treatment Effects by Instructor Type

		Exam Score	Course Grade	Complete Course
Adjunct	Instructor Incentives	0.261***	0.186***	0.067***
		(0.073)	(0.062)	(0.025)
	Combined Incentives	0.119	0.031	-0.034
		(0.093)	(0.091)	(0.034)
	Student Incentives	0.100	-0.087	0.003
		(0.102)	(0.058)	(0.030)
Full-time	Instructor Incentives	0.128	0.002	0.002
		(0.082)	(0.084)	(0.028)
	Combined Incentives	0.180	0.113	0.014
		(0.119)	(0.114)	(0.042)
	Student Incentives	0.020	-0.010	-0.010
		(0.091)	(0.081)	(0.032)
Control group mean	Adjunct	-0.092	-0.026	0.772
		(1.01)	(1.00)	(0.42)
	Full-time	-0.044	-0.015	0.793
		(0.98)	(0.98)	(0.41)
Instructor Incentives:	Pr(Adjunct = Full-time)	0.226	0.075	0.080
Instructors		127	134	134
Observations		5839	6066	6066

Column 1-2: Random effects linear estimation.
Column 3: Marginal effects from random effects probit estimation.
Standard errors in parentheses clustered at the instructor level.
Standard deviation reported for control group mean.
Includes instructor type and instructor type interacted with each treatment.
Dependent variable column 1 and 2: standardized within dept. (mean 0, s.d. 1).
Includes semester and department fixed effects and covariates for student (age, gender, race/ethnicity, baseline credits), instructor (contract value, discount rate) and course (co-requisite).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6: Dosage and Substitution within Program

	Exam Scores	
Panel A: Treatment Effects by Instructor Dosage		
Instructors teaching one course	0.153**	
	(0.073)	$p = 0.588$
Instructors teaching multiple courses	0.202***	
	(0.064)	
Panel B: Treatment Effects by Student Dosage		
Students with one treated course	0.211***	
	(0.064)	$p = 0.446$
Students with multiple treated courses	0.290***	
	(0.102)	
Panel C: Treatment Effects by Student Substitution		
Students with courses in both treatment and control	0.223***	
	(0.060)	$p = 0.191$
Students with all courses in the same treatment	0.340***	
	(0.098)	
Clusters (Instructors)	127	
Total Observations	5839	

Random effects linear estimation. Standard errors clustered at the instructor level.
Last column reports p-value from test of equality between treatment effects in each panel.
Panel A: Observations divided by if the instructor taught multiple courses in a semester.
Panel B: Observations divided by if the student experienced multiple *treatment* courses in a semester.
Panel C: Observations divided by if the student experienced both treatment and control courses in a semester.
The treatment coefficient then compares treatment courses to control courses for students in each of these groups.
Panel A includes controls for total non-program teaching hours.
Panels B & C include controls for the total number of courses a student takes in a semester.
Panel B: The omitted group is never treated.
Includes semester and department fixed effects and covariates for student (age, gender, race/ethnicity, baseline credits), instructor (type, contract value, discount rate) and course (co-requisite).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 7: Spillovers outside Program

	Treatment Semester				Post-Treatment
	Non-Program		Cumulative		Enrollment, Transfer,
	Grades	Completion	GPA	Credits	or Graduation
Instructor Incentives	0.043	0.030***	0.035	0.072	0.004
	(0.036)	(0.012)	(0.044)	(0.227)	(0.012)
Combined Incentives	0.046	0.034**	0.081	0.303	-0.020
	(0.058)	(0.017)	(0.062)	(0.423)	(0.019)
Student Incentives	0.015	0.005	-0.079	-0.402	-0.004
	(0.060)	(0.022)	(0.056)	(0.296)	(0.025)
Control Group Mean	-0.020	0.769	2.455	22.223	0.784
	(0.99)	(0.37)	(1.14)	(19.46)	(0.41)
Instructors	134	134	134	134	134
Observations	5872	5872	6066	6066	6066

Columns 1, 3 and 4: Random effects linear estimation.
Columns 1, 3 and 4: Standard errors clustered by instructor in parentheses.
Columns 2 and 5: Marginal effects from random effects probit estimation.
Columns 2 and 5: Heteroskedasticity-robust standard errors in parentheses.
Column 5: Dependent variable is indicator equal to one if student earns degree, transfers or enrolls in post-treatment semester (Spring 2017 for Fall 2016; Fall 2017 for Spring 2017).
Standard deviation reported for control group mean.
Includes semester and department fixed effects and covariates for student (age, gender, race/ethnicity, baseline credits), instructor (type, contract value, discount rate) and course (co-requisite).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 8: Self-reported expenditures and time use

	Control	Instructor Incentives
Panel A: Expenditure of personal funds (dollars)		
Class materials	66.375	75.863
	(16.031)	(11.232)
Professional Development	96.550	132.182
	(22.129)	(21.441)
Panel B: Time use (hours)		
Teaching class	9.727	9.483
	(0.668)	(0.490)
Preparing for class	5.273	5.944
	(0.467)	(0.442)
Preparing assignments & exams	4.268	4.079
	(0.402)	(0.358)
Grading assignments & exams	5.571	5.586
	(0.516)	(0.426)
Holding office hours	5.939	6.212
	(0.717)	(0.672)
Helping students outside of office hours	3.120	2.568
	(0.392)	(0.282)
Advising students	6.115	5.118
	(0.831)	(0.688)
Administrative work	3.264	3.977*
	(0.438)	(0.412)
Professional development	2.426	2.389
	(0.436)	(0.342)
Outside employment†	8.818	11.132
	(1.268)	(0.757)
Observations	**80**	**97**

Table reports means for each outcome by treatment group.
Standard errors in parentheses clustered at the instructor level.
Observations are at the instructor-semester level.
Significance tests conducted using random effects regression.
† 51 of 101 responses are top-coded for outside employment (> 16 hours).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 9: Treatment Effects on Preference for Loss-Framed Contracts

	Baseline	Fall Endline Preference		
Fall Instructor Incentives	0.073	4.326**	4.529**	4.351***
	(2.307)	(2.069)	(1.782)	(1.666)
Full-Time	-0.191	-0.065	-0.978	0.337
	(2.343)	(1.955)	(1.667)	(1.633)
Constant	-6.134***	-6.055***	-3.579**	-4.636***
	(2.060)	(1.884)	(1.418)	(1.434)
Group 1: (< -$10)			-6.492***	-3.563*
			(2.022)	(1.894)
Group 2: (-$10 to -$0.5)			-8.056**	-6.979**
			(3.417)	(3.485)
Group 4: ($0 to $10)			2.897	-1.960
			(1.967)	(2.429)
Group 5: (> $10)			6.115	0.976
			(10.180)	(8.285)
Discount Rate (δ)				-4.487***
				(1.305)
Pr(Treatment Group Value = $0)	0.001	0.291	0.499	0.850
Instructors	89	90	89	88

Dependent variable: per-student bonus willing to pay for loss-framed contract.
Generalized tobit regression to correct for interval-censored data.
Heteroskedasticity-robust standard errors.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 10: Well-being and stress

	Control	Instructor Incentives
Panel A: Post-Treatment Personal Well-being		
Please rate the overall well-being of you and your family	8.345	8.067
	(0.195)	(0.133)
Please rate your financial security	6.534	6.256
	(0.320)	(0.286)
Do you feel that you have enough time and money for the things that are most important to you?	6.431	5.844
	(0.345)	(0.298)
Do you feel that you are making a difference actively contributing to the well-being of other people and making the world a better place?	8.483	8.055*
	(0.154)	(0.164)
How often do you feel completely stress-free?	5.276	4.689
	(0.355)	(0.285)
How would you usually rate your overall mental health?	8.276	8.253
	(0.203)	(0.152)
Panel B: Post-Treatment Professional Well-being		
Please rate the overall quality of your experience at work	7.483	7.747
	(0.279)	(0.200)
The stress and disappointment involved in teaching at Ivy Tech aren't really worth it	3.759	3.078**
	(0.301)	(0.262)
If I could get a higher paying job, I'd leave teaching as soon as possible	4.017	4.022
	(0.377)	(0.347)
I think about transferring to another school	3.983	3.544
	(0.401)	(0.337)
I don't seem to have as much enthusiasm now as I did when I began teaching	3.483	3.800
	(0.392)	(0.310)
I really enjoy my present teaching job	7.897	7.923
	(0.204)	(0.203)
I am certain that I am making a difference in the lives of the students I teach	8.224	7.747
	(0.182)	(0.207)
If I could start over, I would choose teaching again as my career	7.707	6.527
	(0.257)	(0.331)
My salary is fair for the work that I do	4.517	4.582
	(0.334)	(0.317)
I am treated as a professional by the community	7.466	7.308
	(0.250)	(0.237)
Observations	**81**	**97**

Table reports means for each outcome by treatment group with standard errors in parentheses clustered at the instructor level.
Observations are at the instructor-semester level.
Significance tests conducted using random effects regression with controls for baseline survey responses, semester, and full-time status.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A Appendix Figures and Tables

Figure A.1: Student Incentives Refrigerator Magnet

Figure A.2: Instructor Randomization by Semester

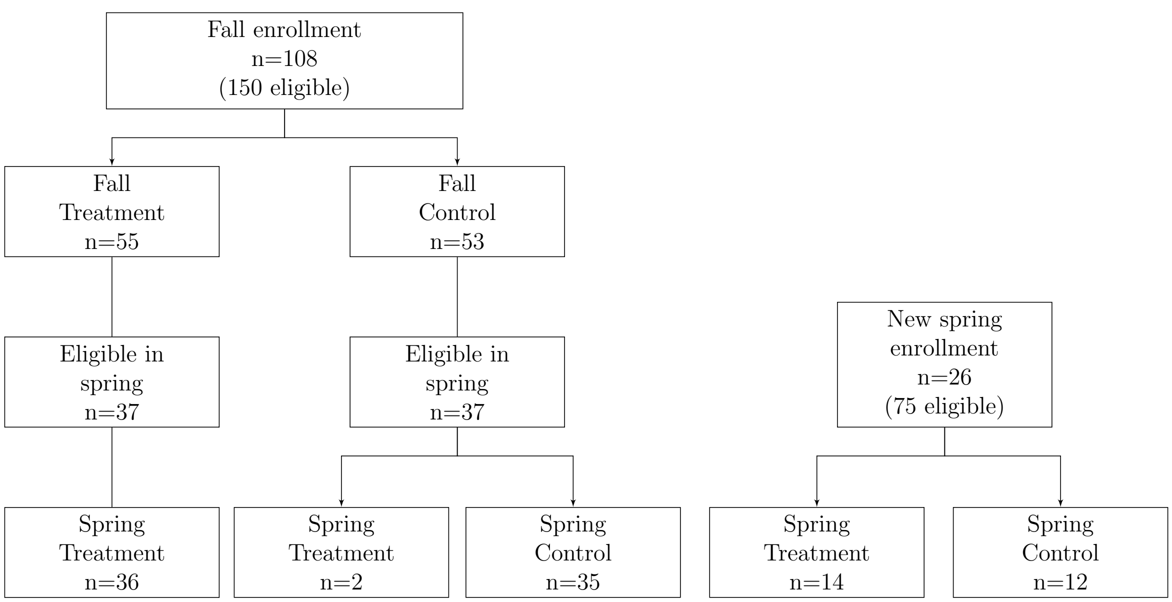

Notes: Arrows indicate random assignment. One Fall Treatment instructor eligible to participate in the spring opted out of continued participation. We randomly assigned two Fall Control instructors eligible in the spring to Spring Treatment. New spring instructors were randomized independently of re-enrolled instructors from the fall.

Figure A.3: Individual Change in Contract Preferences

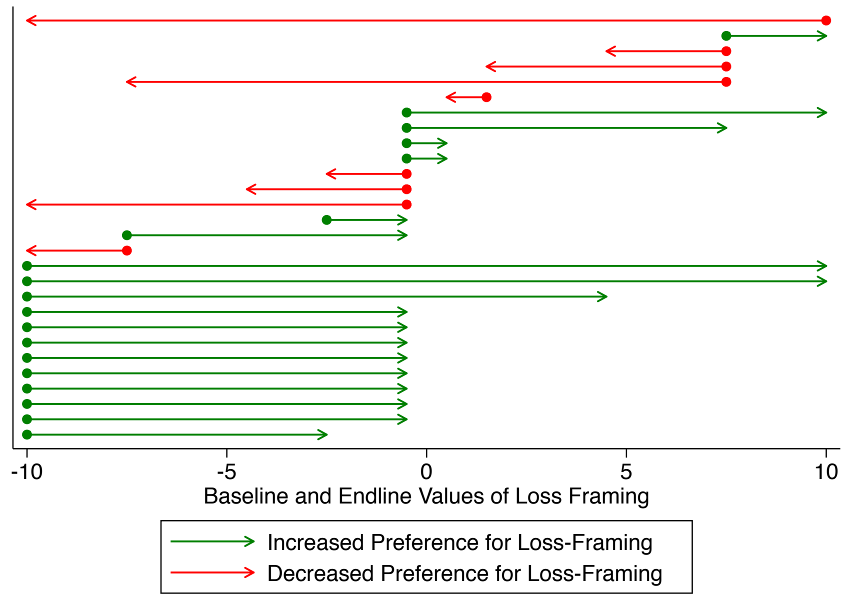

(a) Treatment Instructors

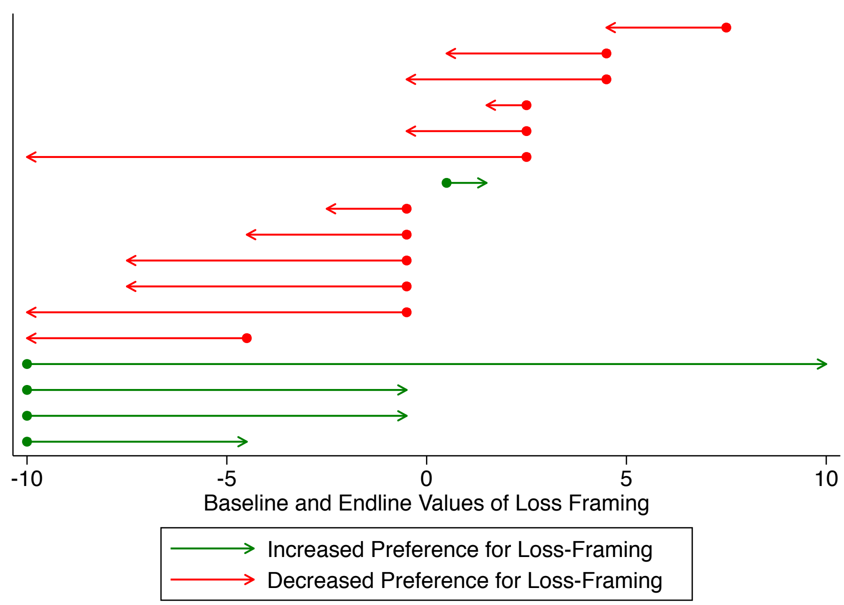

(b) Control Instructors

Notes: Each line represents an individual respondent. Circles indicate baseline preferences and arrowheads indicate endline preferences. Individuals who increased preferences for loss-framing are colored in green, while individuals who decreased preferences for loss-framing are colored in red. Individuals with no change are not represented.

Table A.1: Sample Sizes by Treatment and Semester

		Sample w/ Valid Exam Data		Sample w/ Course Data	
	Treatment	Fall	Spring	Fall	Spring
Observations	Instructor Incentives	1507	624	1592	657
	Combined Incentives	—	687	—	687
	Student Incentives	—	311	—	311
	Instructor Control	1682	1030	1776	1043
	Total	**3189**	**2652**	**3368**	**2698**
Sections	Instructor Incentives	91	42	95	44
	Combined Incentives	—	45	—	45
	Student Incentives	—	23	—	23
	Control	99	70	105	71
	Total	**190**	**180**	**200**	**183**
Instructors	Instructor Incentives	52	50	55	52
	Instructor Control	48	46	53	47
	Total	**100**	**96**	**108**	**99**

Table A.2: Missing Data by Treatment and Semester

	Fall			Spring				
	Control	Instructor Incentives	**Fall F-test**	Control	Instructor Incentives	Combined Incentives	Student Incentives	**Spring F-test**
Missing Demographics	0.021	0.036	**0.206**	0.033	0.025	0.023	0.022	**0.548**
	(0.003)	(0.011)		(0.006)	(0.008)	(0.005)	(0.009)	
Missing Exam Scores	0.055	0.065	**0.731**	0.014	0.049	0.000	0.000	**0.171**
	(0.021)	(0.022)		(0.014)	(0.031)	(0.000)	(0.000)	
Total Missing Data	0.073	0.088	**0.623**	0.047	0.074	0.023	0.022	**0.179**
	(0.020)	(0.022)		(0.015)	(0.031)	(0.005)	(0.009)	
Observations	1815	1652		1079	674	703	318	

Table reports means by treatment group with standard errors in parentheses clustered at the level of randomization (instructor).
"Total Missing Data"—the student is not included in *any* analysis of exams due to missing scores or missing demographics.
p-values are reported for joint orthogonality test across treatment groups within a semester.
Significance tests against control mean: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.3: Effects of Incentives on Exam Scores: Sensitivity Analysis

	OLS Estimation			No Co-Req	Baseline GPA Covariate		Spring
	Pooled	Fall	Spring	Covariate	All	Has GPA	Returners
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Instructor Incentives	0.162***	0.115*	0.223**	0.202***	0.210***	0.173***	0.155*
	(0.060)	(0.061)	(0.092)	(0.055)	(0.055)	(0.062)	(0.088)
Combined Incentives	0.091		0.141	0.145*	0.147**	0.104	0.049
	(0.083)		(0.093)	(0.075)	(0.074)	(0.090)	(0.114)
Student Incentives	-0.051		0.008	0.064	0.086	0.087	0.091
	(0.092)		(0.088)	(0.076)	(0.077)	(0.104)	(0.087)
Department	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Instructor Type	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Baseline Characteristics	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls for Co-Req	Yes	Yes	Yes	No	Yes	Yes	Yes
Baseline GPA	No	No	No	No	Yes	Yes	No
Pr(Instructor = Combined)	0.308		0.236	0.416	0.372	0.368	0.296
Instructors	127	100	96	127	127	127	71
Observations	5839	3189	2650	5839	5839	3393	1953

Dependent variable: exam score standardized within dept. (mean 0, SD 1).

Columns 1-3 use an OLS estimation. Standard errors clustered by instructor.

Columns 4-6 use a random-effects linear estimation. Standard errors clustered by instructor.

Column 4 does not include controls for Co-Requisite courses.

Column 5 controls for baseline GPA. We impute missing GPA as the mean GPA and include an indicator if missing baseline GPA.

Column 6 excludes students missing baseline GPA.

Column 7 only analyzes spring instructors who continued their participation from the fall.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.4: Effects on Exam Scores: Alternative Specifications of Combined Incentives

	Pool Instructor and Combined		Interaction Specification	
	(1)	(2)	(3)	(4)
Instructor Incentives	0.185***	0.192***	0.204***	0.247***
	(0.052)	(0.053)	(0.056)	(0.085)
Student Incentives		0.073	0.065	0.067
		(0.076)	(0.075)	(0.083)
Instructor × Student Incentives			-0.121	-0.148
			(0.101)	(0.110)
Department	Yes	Yes	Yes	Yes
Instructor Type	Yes	Yes	Yes	Yes
Baseline Characteristics	Yes	Yes	Yes	Yes
Instructors	127	127	127	96
Observations	5839	5839	5839	2650

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Random effects linear estimation.

Standard errors in parentheses clustered at the instructor level.

Column (1) pools Instructor Incentives with Combined Incentives and pools Student Incentives with Control.

Column (2) pools Instructor Incentives with Combined Incentives and includes an indicator variable for Student Incentives only. Control is the omitted group.

Column (3) replicates our standard specification using an interaction term in place of a separate indicator for Combined Incentives.

Column (4) repeats Column (3) but restricts the sample to the spring semester.

Table A.5: Student Course Evaluations by Treatment Group

	Control Group	Instructor Incentives	Combined Incentives	Student Incentives
I would recommend this instructor to others	4.366 (0.039)	4.293 (0.048)	4.355 (0.080)	4.355 (0.125)
I would recommend this course to others	4.365 (0.036)	4.329 (0.044)	4.337 (0.073)	4.338 (0.119)
Instructor presents course materials in a well-organized manner	4.461 (0.034)	4.400 (0.040)	4.474 (0.066)	4.416 (0.108)
Instructor presents concepts clearly	4.411 (0.036)	4.273 (0.045)	4.416 (0.071)	4.325 (0.126)
Contents of the assignments contribute to my understanding of the subject	4.495 (0.031)	4.410 (0.040)	4.432 (0.068)	4.455 (0.109)
Instructor is prepared for course	4.563 (0.030)	4.483 (0.039)	4.571 (0.057)	4.487 (0.095)
Instructor is available outside of class (office hours, via e-mail, et cetera)	4.469 (0.032)	4.426 (0.036)	4.456 (0.062)	4.500 (0.100)
Total Responses	752	642	201	82
Response Rate	26.5%	29.6%	30.5%	29.7%

Means reported for each treatment group.

Standard errors clustered at the instructor level.

Significance tests performed using a random-effects linear estimation.

Estimation includes semester and department fixed effects and covariates for student (age, gender, race), instructor (type, contract value, discount rate) and course (co-requisite).

Control group estimated with constant, treatment groups estimated with sum of constant and treatment coefficient.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B Instructor survey: Program description and preference elicitation (for online publication only)

On this page, we describe the **Details of the Bonus Program**

Broadly speaking, you will not have to change anything about the way you teach. We are simply presenting you with the opportunity to earn bonuses based on the performance of your students.

We will randomly select about 60 of the Ivy Tech instructors who sign up for the program to receive bonuses.

Your bonus is based on the number of students you teach that *"Pass the Final"* meaning they receive an A, B or C on a standardized final exam.

Your bonus will be calculated as follows: **you will get $50 per student who passes the final.** For example, if you teach 20 students, you can receive up to $1,000 if all of your students pass the final.

There are two types of bonuses you can receive: **End Bonus** or **Advance Bonus**.

End Bonus:

The End Bonus is simple. You will receive your entire bonus at the end of the semester.

Advance Bonus:

You will receive an upfront bonus at the beginning of the semester equal to the bonus you would receive if half of your students pass the final. At the end of the semester, you will receive or return the difference between your total reward and your upfront bonus.

For example, if you teach 20 students you will receive a check at the start of the semester as if 10 of them pass the final. Your upfront bonus will be $500 ($50 x 10 =$500). If exactly 10 students pass the final, you will keep the Upfront Bonus. If 6 students pass the final, your Total Bonus will be $300 50 x 6 students), and you will have to pay back $200 of your $500 Upfront Bonus. If 14 students pass the final, your Total Bonus will be $700 ($50 x 14 students), and you will receive $200 in addition to your $500 Upfront Bonus.

The bonus will not affect your regular pay from Ivy Tech (your bonus will be paid to you from the University of Arkansas).

We are now going to ask a series of questions about which kind of bonus you prefer to receive from the program: the Advanced Bonus or End Bonus (both of which are described above).

To keep it simple, we will use the example of teaching 20 students. Your actual bonus will be adjusted based on the number of students you teach.

Which would you prefer?

Advance Bonus: $50 Per Student Total Bonus up to $1,000 Upfront Bonus of $500	End Bonus: $50 Per Student Total Bonus up to $1,000 All paid at the end of the semester

---

The next questions ask if you would prefer $50 per student as an Advance Bonus or different amounts of money paid as an End Bonus. Please tell us separately for each question whether you prefer to receive the Advance Bonus or End Bonus. Think of each question independently and simply consider which form of bonus you would prefer to receive.

Please consider the choices carefully because your decisions below will determine the bonus type and amount you receive in the program. We will randomly select one of about 60 instructors in the Bonus Program to receive their choice of the type of bonus they would like to receive. If you are selected to receive the bonus of your choice, we will randomly choose one of your decisions below to be the bonus you actually receive.

For example, if the first question below is chosen to determine your bonus and you tell us that you prefer the Advance Bonus, you will receive $50 per successful student with a $500 upfront bonus. If you tell us you prefer the End Bonus, you will receive $60 per successful student, all paid at the end of the semester.

Which would you prefer?

Advance Bonus: $50 Per Student Total Bonus up to $1,000 Upfront Bonus of $500	End Bonus: $60 Per Student Total Bonus up to $1,200 All paid at the end of the semester

---

Which would you prefer?

Advance Bonus: $50 Per Student Total Bonus up to $1,000 Upfront Bonus of $500	End Bonus: $55 Per Student Total Bonus up to $1,100 All paid at the end of the semester

---

Which would you prefer?

Advance Bonus: $50 Per Student Total Bonus up to $1,000 Upfront Bonus of $500	End Bonus: $54 Per Student Total Bonus up to $1,080 All paid at the end of the semester

---

Which would you prefer?

Advance Bonus: $50 Per Student Total Bonus up to $1,000 Upfront Bonus of $500	End Bonus: $53 Per Student Total Bonus up to $1,060 All paid at the end of the semester

---

Which would you prefer?

Advance Bonus: $50 Per Student Total Bonus up to $1,000 Upfront Bonus of $500	End Bonus: $52 Per Student Total Bonus up to $1,040 All paid at the end of the semester

Which would you prefer?

Advance Bonus: $50 Per Student Total Bonus up to $1,000 Upfront Bonus of $500	End Bonus: $51 Per Student Total Bonus up to $1,020 All paid at the end of the semester

Which would you prefer?

Advance Bonus: $50 Per Student Total Bonus up to $1,000 Upfront Bonus of $500	End Bonus: $50 Per Student Total Bonus up to $1,000 All paid at the end of the semester

Which would you prefer?

Advance Bonus: $50 Per Student Total Bonus up to $1,000 Upfront Bonus of $500	End Bonus: $49 Per Student Total Bonus up to $980 All paid at the end of the semester

Which would you prefer?

Advance Bonus: $50 Per Student Total Bonus up to $1,000 Upfront Bonus of $500	End Bonus: $48 Per Student Total Bonus up to $960 All paid at the end of the semester

Which would you prefer?

Advance Bonus: $50 Per Student Total Bonus up to $1,000 Upfront Bonus of $500	End Bonus: $47 Per Student Total Bonus up to $940 All paid at the end of the semester

Which would you prefer?

Advance Bonus: $50 Per Student Total Bonus up to $1,000 Upfront Bonus of $500	End Bonus: $46 Per Student Total Bonus up to $920 All paid at the end of the semester

---

Which would you prefer?

Advance Bonus: $50 Per Student Total Bonus up to $1,000 Upfront Bonus of $500	End Bonus: $45 Per Student Total Bonus up to $900 All paid at the end of the semester

---

Which would you prefer?

Advance Bonus: $50 Per Student Total Bonus up to $1,000 Upfront Bonus of $500	End Bonus: $40 Per Student Total Bonus up to $800 All paid at the end of the semester