Academic Cheating and Workload Intensity: Theory and Simulations

Richard Marmorstein

Econ 399 - Dr. Guse

Winter 2014

Introduction

Cheating and Honor Codes

Academic cheating is an issue that must be addressed by every university. The issue recently received widespread coverage in the popular press in the wake of the Harvard cheating scandal of 2012. In a class of 279 students, a professor noticed that 125—about half—of the take-home exams looked suspiciously similar. After an investigation, roughly 70 students, more than 1% of Harvard's undergraduate student body, were found guilty of some degree of cheating and were forced to temporarily withdraw from the university (Prez-pea, 2013). This initiated a discussion about the prevalance and severity of cheating in America's universities, but also about possible measures to address cheating, in particular, the adoption of university honor codes.

Honor codes are characterized in academic literature generally as systems which feature written pledges that work is honest, student-run judiciaries that try cheating cases, unproctored examinations, and an obligation of students to report cheating they witness (Melendez, 1985, cited in McCabe et al. (1999)). Honor codes vary in the power the student-run judiciaries hold to punish those found guilty of academic dishonesty.

In this project, I investigate a particular model of punishment featured by some honor codes — a single sanction. Under a single-sanction honor code, the student judiciary must expel any student found guilty of academic dishonesty. The judiciary has no power to institute a lesser punishment, even for minor infractions. Universities whose honor code follows this model include the University of Virginia, the Virginia Military Institution, and Washington and Lee university. The single sanction has been controversial. Recently at the University of Virginia, students voted on a measure to replace the single sanction with a multiple-tiered system, which was rejected (Vasireddy, 2009).

Past research has investigated a number of questions regarding honor codes and academic cheating. To what degree are honor codes effective in preventing cheating? Does this effect vary depending on characteristics of the institution, such as size or selectivity? Is an honor code more likely to prevent cheating by some types of students than others? I was unable to find any study examining specifically the effect of a single sanction. However, McCabe and Trevino (1993) survey students to report how severe they perceive their punishment will be if they are caught cheating–a related factor.

There is a general lack of effective data to conduct empirical research into these questions. Most studies are limited by the available data to investigating the prevalence of cheating on a single campus, such as Vandehey et al. (2007). McCabe and Trevino (1993) use self-reported cheating data from students at multiple universities. They find a significantly higher incidence of cheating among non-honor code universities than among universities with an honor code. Yet, they also find that cheating can be better predicted using other contextual factors, such as how their peers perceive cheating, how likely they perceive it is that they will be reported if they cheat, and how severe the penalty is for cheating. Their approach is not perfect, however. Universities with an established honor code are not a random sample, and likely differ in other characteristics than simply the presence of an honor code. Furthermore, the group of students which choose to fill out the surveys and return them are likely not a random sample of students either, which leaves open the possibility of selection bias. Without better natural experiments, and more detailed and reliable data, these concerns would be difficult to address.

Deviance Theory

Michaels and Miethe (1989) provide an overview of psychological theories of deviance, which are frequently applied to other sorts of illicit behavior such as crime, particularly corporate fraud and white-collar crime. They hold that these theories can also aid in explaining cheating in a university environment. They classify theories of deviance into four groups: deterrence theory, rational choice theory, social bond theory, and social learning theory. Deterrence theory predicts that cheating will be inversely related to the perceived probability of punishment times the severity of the punishment. Rational choice theory extends deterrence theory and additionally predicts that cheaters not only take into account punishment, but also will choose to cheat more as the perceived benefit from cheating increases. Social bond theory predicts that deviant behavior, cheating, will decrease the more 'connected' a community is with regard to attachment, commitment, involvement, and belief. Social learning theory is an adaptation of social bond theory that allows for the possible existence of deviant groups within a community. Among deviant groups, attachment, commitment, etc. will serve to reinforce, and not to diminish, behavior which may be considered 'deviant' for the community as a whole. Each of these theories were found to have degrees of predictive power on the results of self-reported cheating data from students at a large state university.

From the perspective of behavioral economics, Mazar et al. (2008) develop a theory of cheating based on experiments which may resemble an academic environment. They reward students from highly selective universities with money for completing puzzles within a time limit. Suspiciously, students report solving more puzzles when they are given the opportunity to cheat. The researchers vary such factors as the presence of honor pledges (which decreases cheating), substituting monetary rewards with "tokens" that can be redeemed for money. Based on their results, they develop a theory of cheating where cheating behavior is manifested by a tension between external rewards like monetary benefit, and internal rewards, such as being able to think of oneself as honest. Their summary states: "This research shows that people behave dishonestly enough to profit but honestly enough to delude themselves of their own integrity."

These psychological theories provide a useful background to evaluate the plausibility of the theoretical model developed in this project.

Theoretical Model

The model presented here is an adaptation of the model developed by Guse and Marmorstein (Summer 2013). I model a university as a game of n players, who represent students.

Move

Each round, each player i chooses a move with three components: cheating $c_i \in [0, 1]$, effort $e_i \in [0, 1]$, and tolerance $t_i \in [0, 1]$. Higher values of c_i reflect more severe forms of cheating. Higher values of e_i reflect higher levels of honest, academic effort. Tolerance t_i is a value to reflect how willing the student is to observe cheating without *tattling*. Note that the word "tolerance" is not used here to describe a innate, constant personality trait. Rather, students choose a different level of tolerance from round to round as they adjust their perceptions of the cost and benefits of tattling on a cheater.

Tattling

Each round, there is a probability each student will *tattle* on each other student. A student who is tattled on is punished according to either a single-sanction or scheme of *proportional punishment*. In order for tattling to occur, the tattler must first observe the cheater, which is a random event. But a student does not always tattle when he observes cheating. There a second stage of randomness—he is more likely to tattle the lower his chosen tolerance, and the greater the severity of the cheating.

Observation

The probability that student i will observe student j cheating is determined by the exogenous *observation distance* from i to j, denoted as $d(i, j)$. The observation distance is the complement of the probability that, given $c_j > 0$, i will observe j cheating. That is, when $d(i, j) = 0$, student i will always observe the cheating of student j. When $d(i, j) = 1$, student

i will never observe student j cheating. Note that the term "distance" is not being used in its mathematical sense—these values need not satisfy the properties of Euclidean distance, or of metric spaces. Particularly, they may be non-symmetric. It could be that i is more likely to observe j cheating than j is to observe i.

Reporting Observed Cheating

If student i observfes student j cheating, i does not always tattle. Student i is only likely to tattle when $c_j > t_i$. The probability that i tattles on j given that i has observed j cheating is given by $F_i(c_j)$, where F is the cumulative distribution function of a beta distribution with a mean of t_i and a standard deviation of σ, which is a exogenous parameter that applies identically to each student. A beta distribution was chosen to bound the probability between 0 and 1. So long as σ is small, F will assume an S-like shape, assuming probabilities close to 0 for values of c_j below t_i, probabilities close to 1 for values of c_j above t_i, and intermediate probabilities only in a small interval surrounding t_i.

Beta distributions are not usually described by their mean and standard deviation. The typical parameters to the Beta distribution are α and β. Described this way, F_i is the c.d.f. of a Beta distribution with parameters α and β such that

$$\alpha = \left(\frac{1 - t_i}{\sigma^2} - \frac{1}{t_i}\right) t_i^2$$

$$\beta = \alpha \left(\frac{1}{t_i} - 1\right).$$

Utility Function

A student i chooses a move each round to maximize the expected value of his utility function U_i. Each U_i is a Cobb-Douglas utility function combining four components: leisure ℓ_i, success s_i, guilt g_i, and indignance I_i. The purpose of these four components is to help the model tell a plausible story about cheating behavior, reflecting the four psychological theories of deviance mentioned in the background section. Leisure and success capture the benefits and drawbacks students experience as a result of decisions to exert honest effort, or to cheat. They embody the aspect of academic honesty described by deterrence theory and rational choice theory. If students are caught cheating and punished, s_i will assume a low value. This will exhibit a deterrent effect on students' willingness to cheat. Yet, if students get away with cheating, s_i will assume a higher level, increasing the cheaters utility, and allowing the cheater to allocate more time to valued leisure, for a given level of success. In accordance with the rational choice model of deviance, this will induce more cheating the higher the potential rewards. Guilt and indignance reflect more the social theories of deviance. Indignance reflects the disutility students face when students whom they observe cheating do not receive a punishment they perceive to be just. Students become indignant both when they feel a cheater has not been punished harshly enough–or escapes punishment altogether–and also when they feel a cheater has been punished too harshly. Indignance incentivizes students to adjust their tolerance level downwards or upwards, contingent on whether they view punishments as too lenient or too harsh.

The utility function is defined

$$U_i(c_i, e_i, t_i) = \beta_i^\ell \cdot \log(\ell_i) + \beta_i^s \cdot \log(s_i) + \beta_i^g \cdot \log(g_i) + \beta_i^I \cdot \log(I_i)$$

The Cobb-Douglas form indicates our agents experience diminishing returns for each of these components. For example, one hour of leisure time is likely to be more precious to a student who gets little free time than to a student who already has most of his time free. The difference between $30,000 and $40,000 in post-graduation income is likely more salient than the difference between $80,000 and $90,000.

Here, β_i^ℓ, β_i^s, β_i^g, and β_i^I, are weights placed on the components of the utility function. Each individual agent can be assigned his own unique set of weights. For instance, a high value to β_i^g but a low value to β_i^s could represent a student with a strong sense of honor, but little ambition for success. Other combinations are possible, and can model a variety of student types.

Components

Here I present the mathematical definitions of each of the four components.

Leisure

Leisure describes the time a student is free to do what he or she likes, and is not occupied by other activities. In this model, there are two activities which require a student's time: schoolwork, and reporting the cheating behavior of other students.

Leisure is defined

$$\ell_i = (1 - e_i) + \rho \cdot (r - \min(R_i, r))$$

where $\rho \in [0, 1]$ is an exogenous *report cost*, R_i denotes a random variable indicating the number of students on whom i tattles in the round, and $r \in \mathbb{N}$ denotes an exogenous *report cap*. A student does not receive additional disutility for reporting students beyond r. It is necessary to place this cap on reporting, so that ensures that ℓ_i will always be a positive number. If there were no cap on the number of students that could be reported, then, since observation and reporting are probabilistic there would be a miniscule chance that player i would observe and tattle upon every single player j with $d(i, j) < 1$, which for some values of e_i cause ℓ_i to be negative, or near zero. Since U_i contains the log of ℓ_i, this would result in $\mathrm{E}(U_i)$ being undefined ($-\infty$.)

Success

Success describes the benefits of putting forth high academic effort, and of unpunished cheating. It also reflects the drawbacks of punished cheating. The success component is also a random variable. Its value depends on whether student i is observed and tattled on, which are random processes.

Each student is assigned an exogenous success endowment $\omega_i^s \in \mathbb{R}^+$. This endowment is received in every case.

If the student is not punished, then

$$s_i = \omega_i^s + c_i + (1 - w_i) \cdot e_i$$

where w_i describes the *workload* faced by the student. Workload is intended to embody the differences between how much effort is required to be successful at different universities. With high workloads, it takes a larger amount of effort to yield the same amount of success.

If the student is punished, then s_i depends on the punishment scheme of the university. Does the university have single sanction honor code or does it follow a *proportional punishment* scheme? Under a single sanction honor system:

$$s_i = \omega_i^s$$

Under proportional punishment

$$s_i = \omega_i^s + (1 - c_i) \cdot (1 - w_i) \cdot e_i$$

Guilt

Guilt is simply disutility from cheating. There is an endowment $\omega_i^g \in \mathbb{R}^+$, which adjusts the sensitivity that a student has to differences in levels of cheating.

$$g_i = \omega_i^g + c_i$$

Because U_i wraps g_i in a log, students with low values of ω_i^g perceive a larger difference between a low c_i and a high c_i.

Indignance

Indignance describes the disutility students experience when they observe injustice. Injustice occurs in three flavors: when a student cheats but is not punished at all, when a student cheats but is punished 'too leniently', and when a student cheats but is punished too harshly. The formula for indignance varies from a single-sanction environment to a proportional punishment environment. In a single-sanction environment the formula is

$$I_i = \beta_I \log(\omega_I + \sum_{j \in \text{Pun}} |1 - max(1, \frac{1}{\gamma_I} c_j)| + \sum_{j \in \text{Free}} |max(1, \frac{1}{\gamma_I} c_j)|)$$

Under a proportional punishment scheme the formula is

$$I_i = \beta_I \log(\omega_I + \sum_{j \in \text{Pun}} |c_j - max(1, \frac{1}{\gamma_I} c_j)| + \sum_{j \in \text{Free}} |max(1, \frac{1}{\gamma_I} c_j)|)$$

where ω_i^I is an endowment of indignance that adjusts sensitivity, PUN indicates the students which student i observed cheating and were punished, and FREE indicates the students who got off scot-free. γ_i is another parameter called *lenience*, which defines a students ideal punishment for those he observes cheating. At γ_i, and all higher levels, the ideal punishment is 100%. At $c_j = 0$, ideal punishment is 0%. Ideal punishment is linear on the intermediate values $[0, \gamma_I]$, defining a function for ideal punishment: $\max\left(1, \frac{1}{\gamma_I} c_j\right)$.

Notice that indignance is more intense the greater the severity of the observed cheating.

There are anecdotes that teachers and students at schools with single-sanctions sometimes forgo turning cheaters in, because they do not wish to impose such severe consequences upon them. This model of indignance captures that behavior. An agent with a high "lenience" will be incentivized to lower their tolerance under a single sanction, to lower the incidence of punishments they generally feel to be too harsh.

Simulation

Because the model described above is so complex, I do not attempt to find and analyze the equilibria of the game. Instead I report and analyze the results of simulations embodying the model.

Learning

An analytical solution to the model would describe a Nash Equilibrium—a set of moves where each player's move was a best response to every other player's move. Analyzing the movements of such solutions under changing parameters is a compelling way to characterize a model, because these equilibria are the only points where no player has an incentive to change, and so seem like likely outcomes. A particular advantage of this analysis is its extreme mathematical rigor. But there are drawbacks to this approach. It frequently requires the formal assumption that each player of the game is perfectly rational, and that the parameters of the game and the payoff functions of the other players are all common knowledge. While analyzing the behavior of perfectly rational agents can be very useful for understanding real human behavior, it is also true that, typically, mortals on Earth do not actually possess perfect rationality and infallible game theoretic insight into the motives of their fellow mortals.

An alternative approach involves modelling games as repeated processes of learning. A game is repeated over several rounds, and every player updates his move each round based on some heuristic drawn from the previous moves of other players. A straightforward technique for learning is *fictitious play*, where players choose their move to maximize expected utility assuming the other players in the game are playing a mixed strategy corresponding to their history of play. Because moves in this project are already continuous, I choose an even simpler learning mechanism: players choose the moves which maximize utility assuming the other players in the game are playing the move which they played the previous round.

These learning techniques do not carry the same mathematical precision as Nash Equilibrium analysis. They are not guaranteed to converge in all circumstances. Any points of convergence are not guaranteed to be unique, and are sensitive to initial conditions. However, a learning approach does reflect how in reality, actors do predict the future based on the past, and adapt their choices as their environment changes.

Simulator

I wrote a simulator in Python to embody the model described above and generate a data set for output. Here I describe how the simulator operates.

Game Parameters

First, the simulator constructs a digital representation of the game embodying a set of *game parameters* which are in turn calculated from a set of *configuration parameters*. Game parameters consist of all the exogenous parameters present in the definition of the game. This includes the number of players n; parameters which apply universally throughout the game, such as the cost of reporting, ρ; and parameters which are unique for individuals such as each β_i coefficient, or lenience γ_i. Observation distances $d(i, j)$ are also game parameters.

Execution Parameters

Execution parameters are values which play a role in simulation but are not present in the definition of the game. They may concern either the simulators learning mechanism–such as each agent's initial move, or the length of the simulation—the number of rounds which the game is repeated. They may also determine details concerning the internals of the simulation software, which are described in more detail later.

Configuration Parameters

I do not determine each game parameter and execution parameter individually. Doing so would greatly complicate analysis, since the number of game and configuration parameters is itself dependent on n. Instead, I determine each game parameter based on a set of *configuration parameters*. After all game and execution parameters have been configured, I use them to execute two simulations—one with a single sanction, one without—but identical in every other aspect.

For example, one configuration parameter I call "β^ℓ mean". This parameter determines the mean of an exponential distribution from which I sample each individual β_i^ℓ game parameter. I choose an exponential distribution for distributing the β_i coefficients because exponential distributions are guaranteed to be positive, and furthermore are the *maximum entropy probability distribution* for the class of distributions with a given mean restricted to the positive domain. This means, in terms of information theory, choosing an exponential distribution imposes the least amount of structure onto the β_i parameters, given those requirements.

Another example of a configuration parameter is *nbrs*, which determines the n^2 observation distances between each pair of agents. In this project, distances are chosen as if the n agents were uniformly arranged in a circle in Euclidean space. It would be possible to use the radius of the circle as a configuration parameter directly, but that has the disadvantage that a given radius might describe a sparsely "packed" group of students in simulations with a small n but a densely packed group with a large n. Instead, I calculate the appropriate radius from "*nbrs*", which determines the number of *neighbors*—other agents which are present within a one unit radius of each agent. Another advantage of this specification is that it would be easily generalized, and could easily apply in arrangement schemes besides circles.

Sampling Configuration Parameters

My approach to characterizing the model involves examining the effect of varying each configuration parameter on which moves are chosen after many rounds of simulation. This requires a dataset with an appropriate variance of the configuration parameters. Towards this end, I manually specify a configuration parameter space. This consists of an upper and lower bound for each parameter. For each simulation, I sample a value for each configuration parameter from the uniform distribution with the specified bounds.

A complete list of each configuration parameter, a description of its relationship to the game or execution parameters, and its lower and upper bound is given in Table 1 on page 17.

Execution

A single simulation is executed as follows. First, I sample a set of configuration parameters from the configuration parameter space. My simulator uses these to select a set of game and execution parameters, as described above. One product of this process is an initial move for each player. What remains is to calculate each player's move for the remaining rounds. Player i's move in round $k > 0$ will be a move which satisfies

$$\max_{(c_i,e_i,t_i)\in[0,1]^3} \mathbb{E}(U_i(c_i, e_i, t_i))$$

which conditional on the other players moves taking on the same value they did in round $k-1$, and also depends on the values of the game parameters selected prior to the simulation.

Optimization

This is a bounded optimization problem of a function of unspecified form, which is not guaranteed to be continuous. For such a problem, precisely locating a global optimum—or even a local optimum—cannot be guaranteed. But algorithms exist to find approximate solutions. My simulator leverages a Python library called NLopt, which provides a common interface to execute any of several algorithms intended for approximating non-linear optimization problems. I have chosen to leverage an evolutionary algorithm known as *Improved Stochastic Ranking Evolution Strategy* . The details of the algorithm are beyond the scope of this paper, but in general, evolutionary optimization strategies operate by guessing points in the domain, evaluating the objective function at those points, and using the information gleaned from previous guesses to produce successively better guesses, favoring areas near guesses which produced high objective function values. This approach is very flexible and can be adapted to a variety of objective functions—but the flexibility comes at a performance cost. Optimization using this method requires the computer to evaluate the expected utility for each agent many, many times.

Approximation

Precisely calculating the value of the expected utility function is computationally expensive. For instance, in a game of n players, there are potentially 2^{n-1} different combinations of students whom an individual agent might observe cheating. The combinations would occur with different probabilities and would be associated with different levels of utility from indignance and leisure, each of which would have to be calculated and summed across to derive the true value of the expected utility function. Because the evolutionary optimization technique already requires a large number of function evaluations, this is computationally infeasible even for low values of n.

To overcome this, I instead approximate the value of the expected utility function. Because

$$\begin{aligned}
\mathbb{E}(U_i(c_i, e_i, t_i)) &= \mathbb{E}\left(\beta_i^\ell \cdot \log(\ell_i) + \beta_i^s \cdot \log(s_i) + \beta_i^g \cdot \log(g_i) + \beta_i^I \cdot \log(I_i)\right) \\
&= \beta_i^\ell \cdot \mathbb{E}[\log(\ell_i)] + \beta_i^s \cdot \mathbb{E}[\log(s_i)] + \beta_i^g \cdot \mathbb{E}[\log(g_i)] + \beta_i^I \cdot \mathbb{E}[\log(I_i)]
\end{aligned}$$

an approximation for expected utility can be readily calculated from approximations of the expectation of the log of each component (this is not necessary for guilt, which is not random). I use a Monte Carlo method to produce my approximations—that is, I use a random number generator and produce a large number of realizations of values for the log of each component and average across those realizations to approximate the expectation. The number of realizations across which to average is an execution parameter: "estimation iterations".

Play

Starting with the round after the initial move, the simulator uses the optimization method above to calculate each agent's selected move. Once it calculates a move for each student, it uses those moves to generate the next round, and so forth, until the specified number of rounds has been reached. The move of every student in every round is output–accompanied by the configuration parameters and a selection of game parameters relevant to that student. I run many experiments this way, and compile the total output into a data set for analysis.

Data

Description

I analyze a dataset produced according to the configuration parameter space described in Table 1. It contains a sample of 1579 pairs of simulations. Descriptive statistics describing the outcome variables are presented in Table 2. These are experiment-level averages and variances across all students for each move parameter. For example, in an average experiment, the average student cheated at a level of 0.208 in the final round. The standard deviation of the average final cheating level of an experiment was 0.172, and the average variance of students' average cheating levels in the final round was 0.092, which implies a standard deviation of about .3.

Table 1: Description of Configuration Parameter Space

Parameter Name	Lower Bound	Upper Bound	Parameter Description
n	5	25	Number of Agents
nbrs	2	10	(Step by 2) Observation distances between each pair of students is set as if the students were evenly spaced around a circle in euclidean space, whose radius is determined so that each student has *nbrs* other students within 1 unit of distance.
β_ℓ mean	0	1	Each agent's β_ℓ parameter was sampled from an exponential distribution with this mean.
β_s mean	0	1	Mean of exponential distribution for β_s
β_g mean	0	1	Mean of exponential distribution for β_g
β_i mean	0	1	Mean of exponential distribution for β_I
lenience mean	0	1	γ_i for each student is randomly sampled from a truncated normal distribution between 0 and 1, with this mean.
lenience sd			The standard deviation of the truncated normal from which γ_i is selected.
workload mean	0	1	w_i for each student is randomly sampled from a truncated normal distribution between 0 and 1, with this mean.
ω_g mean/sd	0	1	Mean and sd of truncated normal distribution for ω_g
ω_s mean/sd	0	1	Mean and sd of truncated normal distribution for ω_s distribution for ω_s
ω_I mean/sd	0	1	Mean and sd of truncated normal distribution for ω_I
rCost	0	0.5	Directly determines ρ, the reporting cost.
rCap	0	5	Directly determines r, the reporting cap.
betasd mean	0.001	0.001	Directly determines σ, the standard deviation of the beta distribution whose cdf determines the likelihood that student i will report student j if he observes his cheating. The lower this number, the more the cdf resembles a discrete jump at student i's tolerance level.
estiters mean	20	20	Expectations for the utility function are not calculated directly, but through realizing random components of the utility function several times and then averaging them. This parameter determines the number of realizations across which to average.
initlow	0	0	Initial move components are sampled from a uniform distribution with this lower bound
inithigh	1	1	Initial move components are sampled from a uniform distribution with this lower bound
n rounds	15	15	The duration of the simulation.

Table 2: Final moves summarized

Statistic	Min	Max	Mean	St. Dev.
c_i mean	0.0001	0.948	0.208	0.172
e_i mean	0.003	0.979	0.329	0.210
t_i mean	0.184	0.986	0.592	0.130
c_i variance	0.000	0.289	0.092	0.069
e_i variance	0.00000	0.291	0.106	0.060
t_i variance	0.00003	0.202	0.074	0.031

The distributions of the outcome variables are visually depicted in the histograms in the left column of Figure . The distributions seem to describe a feasible set of academic institutions. As you might expect, the average level of cheating is skewed heavily toward zero. Institutions where the average cheating level is above 0.5 are extremely rare. The distribution of effort also is feasible. Most institutions experience a low average level of effort, but high-effort institutions are definitely present. The distribution of tolerance indicates that there are no institutions where the average student is committed to reporting all low levels of cheating, and very few institutions where the average student is hesitant to report any cheating at all. My own guess is that in reality, the average student at most universities are more tolerant of cheating than this. This might not be all unfortunate, though. Because reporting by professors or other non-student agents is not captured by this model, what may be an unrealistic distribution of tolerance *coming from* students may be a more realistic of the distribution of tolerance *faced by* students.

Figure depicts the average level of each move component over time. In round 0, each move component averages near 0.5, since initial moves are randomly chosen. Quickly, the curves adjust and level out. This is reassuring, and indicates that these games are not random or chaotic systems.

Figure 1: Distribution of Final Average Moves

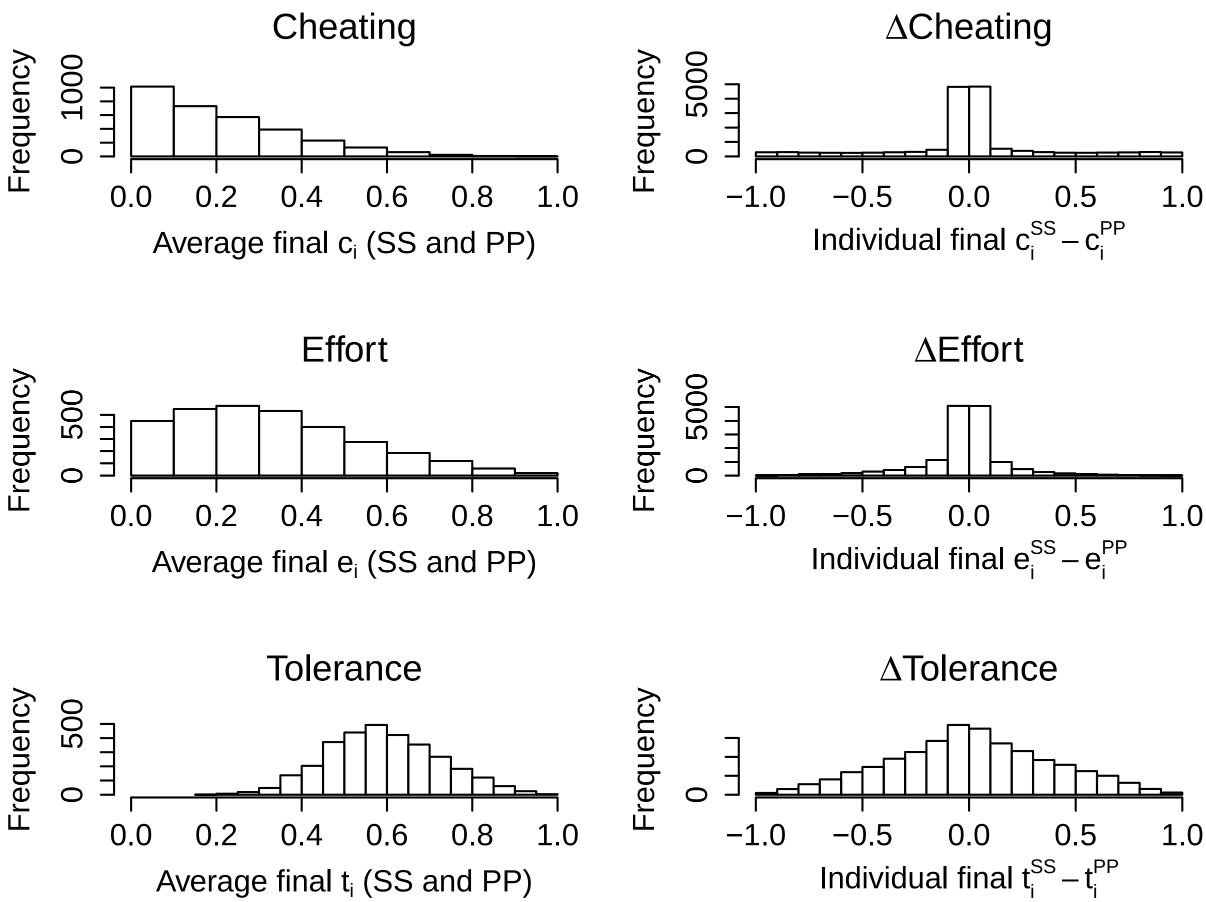

Noteworthy Results

Table 3 captures the main story driving the model. It displays the results of plain, OLS regressions at the experiment level. The explanatory variables are the configuration parameters used to generate the experiment, and the dependent variables are the average levels of the move components for that experiment in the final round. The first four columns serve as a sanity test to verify that the basic elements of the model are operating as they ought to. Putting more weight on leisure leads to less effort, and greater tolerance, because schoolwork and tattling are both hassles. More weight on success leads to greater cheating and effort because those lead to success. Putting more weight on guilt decreases cheating–but also tolerance. An explanation for that result is less obvious. A somewhat surprising observation

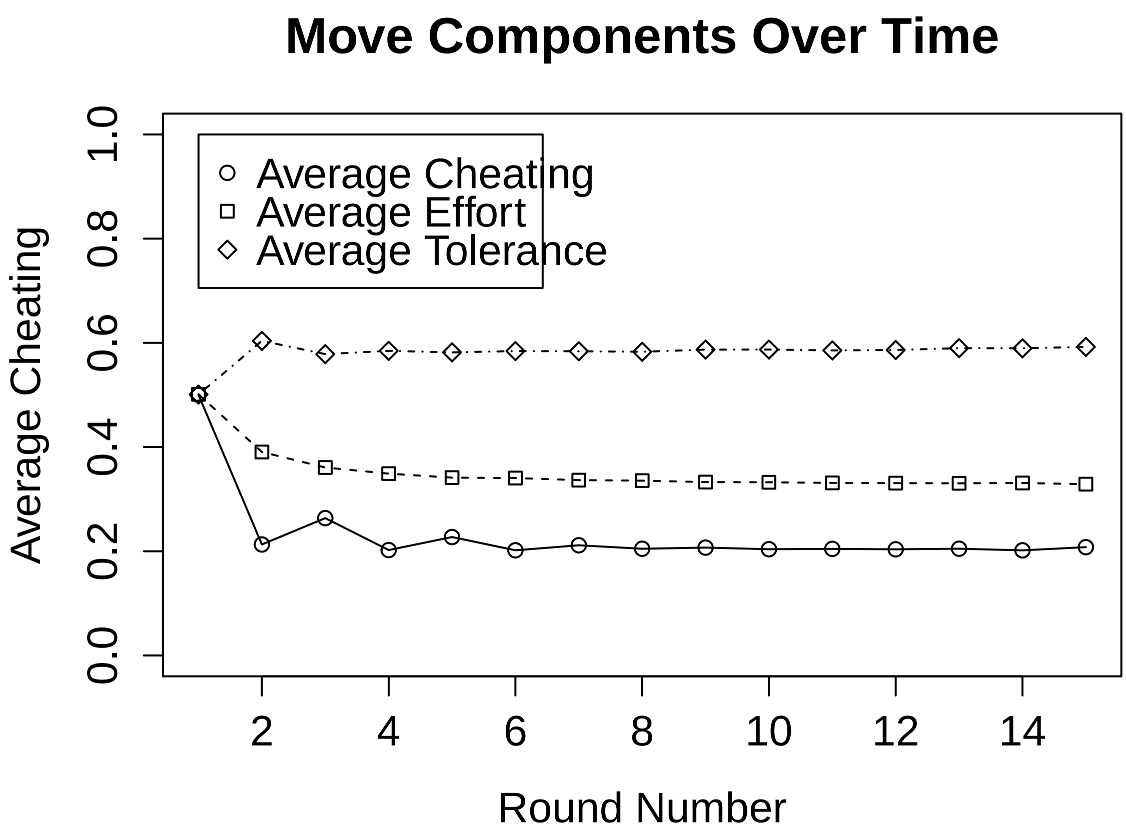

is that neither workload nor lenience seem to have a measured effect on any of the result variables. Finally, the primary result of interest is found at the bottom of the table. A single sanction appears to have no effect on cheating or tolerance, but effort declines slightly in the presence of a single sanction.

It is difficult to understand how to interpret magnitudes in this context. There is no such thing as a "cheating unit," and it is impossible to attach any of these values to concrete events in reality. Clearly, the most influential factors in the model are the β coefficients. The coefficient in the table refers to the effect of a unit change in one of these variables. A unit change, however, is the entire range of these beta coefficients. For example, an institution where the average student attaches no weight at all to conscience is expected to have an average cheating level 0.281 higher than an institution where the average student attaches the maximum weight. The standard deviation of the beta coefficients is about 0.29,

Table 3: Average experiment-level regressed on average final move components

	Average final value		
	c_i	e_i	t_i
	(1)	(2)	(3)
β_ℓ mean	0.014*	−0.396***	0.076***
	(0.008)	(0.008)	(0.007)
β_s mean	0.208***	0.319***	0.014*
	(0.008)	(0.008)	(0.007)
β_g mean	−0.283***	0.005	−0.094***
	(0.008)	(0.008)	(0.007)
β_i mean	−0.007	−0.018**	−0.010
	(0.008)	(0.008)	(0.007)
ω_s mean	−0.002	0.007	−0.004
	(0.008)	(0.008)	(0.007)
ω_g mean	0.142***	0.010	0.054***
	(0.008)	(0.008)	(0.007)
ω_i mean	−0.002	0.006	0.010
	(0.008)	(0.009)	(0.008)
γ_i mean	0.00001	−0.006	−0.005
	(0.008)	(0.008)	(0.007)
γ_i sd	0.019**	0.007	−0.006
	(0.008)	(0.008)	(0.007)
w_i mean	0.003	0.007	−0.001
	(0.008)	(0.008)	(0.007)
r	0.005***	0.035***	0.023***
	(0.002)	(0.002)	(0.002)
ρ	0.029*	0.246***	0.128***
	(0.016)	(0.017)	(0.015)
Single Sanction	0.003	−0.019***	0.007*
	(0.005)	(0.005)	(0.004)
Observations	3,158	3,158	3,158
R^2	0.424	0.589	0.171
Note:			*p<0.1; **p<0.05; ***p<0.01

the standard deviation of the standard uniform distribution. So, a reasonable change in β^s would result in a 0.29 * 0.285 = 0.083 increase in effort, a bit more than three times the decrease in effort associated with the introduction of a single sanction. So, imagine a university where the average student is about a standard deviation lazier than the average student at a second university. Picture the difference in the amount of effort they put in to their schoolwork. Maybe the student at the lazier university does 3 hours less of studying a week. I just made that number up. Then, under this interpretation, under a single sanction the average student studies 3 / 3 = 1 hour less a week. I should emphasize that I do not take sort of magnitude analysis very seriously. My goal in producing that number is simply to argue that the coefficient on the single sanction is large enough to be worth analyzing–I am not trying to use the model to make any prediction about the actual effects of a single sanction.

In Table 4 I present results from the analysis of a dataset including the individual agents in all the experiments. On the left hand side is the level of each move component chosen by the agent in the final round. On the right hand side are that agent's β_i coefficients, ω_i endowments, γ_i, and w_i. These regressions also include experiment fixed effects. Thus, the coefficient in the first column on β^ℓ describes the effect on an individual's cheating level within an experiment. The regression described by the rightmost three columns include on the right hand side, the final value of the two move components which are not on the left hand side.

Many of these results coincide with the experiment-level findings described in Table 3. Liking leisure decreases effort. Liking success increases effort and cheating. Having a conscience decreases cheating. There are some important discrepancies, however. Whereas at the experiment level, increasing the average workload in an experiment had no effect on the average level of final effort, increasing the average workload of an individual student does increase his final level of effort as you would expect based on the specification of the model, which is reassuring. This raises the question, though, why is this effect not registered at the

Table 4: Individual-level Final Moves, with Experiment-level Fixed Effects.

	c_i	e_i	t_i	c_i	e_i	t_i
β^ℓ	0.02^{***}	-0.18^{***}	0.02^{***}	0.02^{***}	-0.18^{***}	0.02^{***}
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
β_i^s	0.13^{***}	0.21^{***}	-0.04^{***}	0.12^{***}	0.21^{***}	-0.02^{***}
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
β_i^g	-0.12^{***}	0.01^{**}	0.01^{**}	-0.12^{***}	0.01^{**}	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
β_i^I	0.00	0.00	-0.01^{***}	0.00	0.00	-0.01^{***}
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
ω_i^s	-0.01^{*}	-0.01	0.00	-0.01^{*}	-0.01	0.00
	(0.01)	(0.00)	(0.01)	(0.01)	(0.00)	(0.01)
ω_i^g	0.09^{***}	0.01	0.00	0.09^{***}	0.01	0.01
	(0.01)	(0.00)	(0.01)	(0.01)	(0.01)	(0.01)
ω_i^I	-0.01	0.00	0.00	-0.01	0.00	0.00
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
workload	0.01	-0.32^{***}	0.00	0.01	-0.32^{***}	-0.01
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
lenience	-0.02^{*}	0.00	0.00	-0.02^{*}	0.00	0.00
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
e_i				0.01		-0.03^{***}
				(0.01)		(0.01)
t_i				-0.06^{***}	-0.03^{***}	
				(0.01)	(0.01)	
c_i					0.01	-0.06^{***}
					(0.01)	(0.01)
R^2	0.31	0.54	0.16	0.31	0.54	0.16
Adj. R^2	0.27	0.51	0.11	0.27	0.51	0.12
Num. obs.	30336	30336	30336	30336	30336	30336

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$

Table 5: Statistical models

experiment level?

In Table 6, I explore which variables determine the size of the difference between a single sanction experiment's result and that of its identical twin proportional punishment experiment. An important result here is the coefficient on β^s for effort, which is even greater than the coefficient in Table 3 for the single sanction. This means that the experiments describing institutions who place the highest premium on success experience an even greater decrease in effort in the presence of a single sanction—but that experiments with a low average weight on success do not experience much of an effect at all. There are also coefficients of some size on the β^s coefficient for tolerance and the β^g coefficient for cheating–however these coefficients are very nearly insignificant. The most influential factor in determining the size of any difference is ρ. It is difficult to interpret ρ because, while ρ increases the cost of reporting–it also has an effect on the size of the endowment of leisure time. In this dataset, ρ varies from 0 to 0.5. So dividing the coefficients in half describes the effect of taking ρ from its minimum to its maximum value. The coefficients indicate that, when reporting a cheater is not much of a hassle, there is less cheating under a single sanction. But when reporting a cheater is a big hassle, there is more cheating under a single sanction. Also, when reporting a cheater is not a hassle, students are more willing to report under a single sanction, but when it is a hassle, they are less willing to report.

It appears that, in addition to not affecting final moves individually, workload does not have an effect on them jointly with a single sanction, either. This result means that there are only the only remaining things to say on the subject of workload is that it does not appear to change the effectiveness or ineffectiveness of a single sanction in this model.

Analysis

Why does a single sanction decrease effort? One explanation is that it makes effort more risky for cheaters, and therefore a less attractive option. Cheating is always risky. Under both systems, if you get caught, the cheating will have done you no good. But under a

Table 6: Determinants of the difference between single-sanction and proportional punishment

	Average final value		
	$c_i^{SS} - c_i^{PP}$	$e_i^{SS} - e_i^{PP}$	$t_i^{SS} - t_i^{PP}$
	(1)	(2)	(3)
β_ℓ mean	0.005	−0.004	0.002
	(0.012)	(0.006)	(0.012)
β_s mean	−0.005	−0.025***	0.026**
	(0.012)	(0.006)	(0.012)
β_g mean	−0.025**	0.006	0.002
	(0.012)	(0.006)	(0.011)
β_i mean	0.003	0.009	−0.017
	(0.012)	(0.006)	(0.012)
ω_s mean	−0.007	0.008	−0.006
	(0.012)	(0.006)	(0.011)
ω_g mean	0.009	−0.007	0.008
	(0.012)	(0.006)	(0.011)
ω_i mean	−0.008	−0.007	0.004
	(0.012)	(0.006)	(0.012)
γ_i mean	0.023**	0.010*	0.015
	(0.012)	(0.006)	(0.011)
γ_i sd	0.028**	−0.004	−0.008
	(0.012)	(0.006)	(0.012)
w_i mean	−0.005	0.003	−0.010
	(0.012)	(0.006)	(0.012)
r	0.002	0.001	0.0002
	(0.002)	(0.001)	(0.002)
ρ	0.071***	−0.010	0.066***
	(0.023)	(0.012)	(0.023)
Constant	−0.027	−0.015	−0.017
	(0.021)	(0.011)	(0.021)
Observations	1,579	1,579	1,579
R^2	0.016	0.020	0.012
Note:	*p<0.1; **p<0.05; ***p<0.01		

single sanction, effort is risky too. No matter how much effort you expend, if you are caught cheating, you are no better off.

If this is the force behind the single sanction's effect, there should be a larger effect where a single sanction renders effort more risky. Effort is more risky the higher the level of cheating by the agent in question, and also the lower the tolerance of the other students at the institution. So to test this hypothesis, I run a regression with effort as the dependent variable on individual-level data of final effort, including both individual and simulation parameters as controls, and including interaction terms between a single sanction and the students effort, and between a single sanction and the average tolerance over all agents in the experiment. The results are reported in Table 7.

Turns out, the data is consistent with this theory. Agents who cheat more see *more* of a decrease in effort under the single sanction, and with significance. The sign of the average t_i is in the opposite direction of what is predicted, but is insignificant.

Table 7: The difference between an individual final effort under a single sanction and his proportional punishment identical twin's final effort

	Individual final value
	final $c_i^S S - c_i^P P$
Simulation Average t_i	0.011
	(0.014)
This agent's c_i (in the SS)	-0.031^{***}
	(0.006)
Constant	-0.019
	(0.013)
β_i^ℓ	Yes
β_i^s	Yes
β_i^g	Yes
β_i^I	Yes
ω_i^s	Yes
ω_i^g	Yes
ω_i^I	Yes
w_i	Yes
γ_i	Yes
β^ℓ mean	Yes
β^s mean	Yes
β^g mean	Yes
β^I mean	Yes
ω^s mean	Yes
ω^g mean	Yes
ω^I mean	Yes
γ mean	Yes
γ sd	Yes
r	Yes
ρ	Yes
n	Yes
Neighbors	Yes
Observations	15,168
R^2	0.009
Adjusted R^2	0.008
Residual Std. Error	0.195 (df = 15143)
F Statistic	5.970^{***} (df = 24; 15143)
Note:	*p<0.1; **p<0.05; ***p<0.01

Conclusion

In this project I have exposited a theoretical model of academic cheating with an emphasis on realism, developed an extensible simulator embodying the model, established a system for describing the parameters of the model, and developed methods for generating and analyzing the data. I analyze one data set from a varied parameter space, and begin to characterize the model through conducting a series exploratory regression analyses. Through these analyses I discovered that, in this model, a single sanction, for the average case in this dataset, does not affect cheating or tolerance, but does cause a slight but important decrease in effort. A single sanction causes different effects in specialized situations. It can influence the average level of cheating upwards or downwards—upwards when the time cost of reporting a cheater is low, and downwards when it is high. In those circumstances, it also influences tolerance in the same direction.

I more deeply explore one explanation for the single sanctions downward pressure on effort—that under this model's single sanction, effort is risker and so delivers less payoff on average, inducing agents to reduce it. In the data, the effect is more pronounced when the riskiness of effort is more intense, which is compatible with this explanation.

This project clears the path for several avenues for potential future research. There are still anomalies in the data which have not been explained. Why does a single sanction affect average cheating and tolerance at the individual level, but not at the experiment level? Projects building on this one might more deeply explore the impact of different social structures on cheating, and determine the observation distances between students more cleverly than placing them in a circle. Perhaps future research could introduce an additional sort of distance beyond "observation distance" which reprented more complicated social ideas such as friendship or respect. Such extensions would be straightforward using the simulator I have written.

References

Mazar, N., O. Amir, and D. Ariely (2008). The dishonesty of honest people: A theory of self-concept maintenance. *Journal of marketing research 45*(6), 633–644.

McCabe, D. L. and L. K. Trevino (1993). Academic dishonesty: Honor codes and other contextual influences. *Journal of Higher Education*, 522–538.

McCabe, D. L., L. K. Trevino, and K. D. Butterfield (1999). Academic integrity in honor code and non-honor code environments: A qualitative investigation. *Journal of Higher Education*, 211–234.

Melendez, B. (1985). *Honor code study.* Harvard University.

Michaels, J. W. and T. D. Miethe (1989). Applying theories of deviance to academic cheating. *Social Science Quarterly 70*(4), 870–85.

Prez-pea, R. (2013, September). Students accused of cheating return awkwardly to a changed harvard. *The New York Times*.

Vandehey, M., G. Diekhoff, and E. LaBeff (2007). College cheating: A twenty-year follow-up and the addition of an honor code. *Journal of College Student Development 48*(4), 468–480.

Vasireddy, P. (2009, February). Students reject multiple-sanction honor referendum. *Cavalier Daily*.

Acknowledgements

- Thank you to my advisor Dr. Joseph Guse for guidance throughout this whole process, and for the big ideas behind the model and analysis technique.

- Thanks to Dr. Katie Shester for valuable feedback during the fall semester.
- Thank you to Dr. Art Goldsmith for the good idea of putting result variables on the right hand side of the regression equations—during the poster session.