Grade Inflation

William Chan
The University of Hong Kong

Li, Hao
University of Toronto

Wing Suen
The University of Hong Kong

October 11, 2002

Abstract: When employers cannot tell whether a school truly has many good students or whether it is just giving easy grades, schools have an incentive to inflate grades to help their mediocre students. However schools also care about preserving the value of good grades for their good students. We construct a signaling model in which grade inflation is the equilibrium outcome. The inability to commit to an honest grading policy in an environment of private information reduces the informativeness of grades and hurts the school. We also show that grade inflation by one school makes it easier for another school to fool the market with grade inflation. Hence easy grades are strategic complements, and this provides a channel to make grade inflation contagious.

1. Introduction

Grade inflation is a persistent problem in higher education. Beginning in the late 1960s, a number of studies have documented an upward trend in college grade point averages (e.g., Juola 1968; Carney et al. 1978; McKenzie 1979; Kolevzon 1981; Millman et al. 1983; Sabot and Wakerman-Linn 1991). Grade inflation has been felt in elite universities in the United States (e.g., *The New York Times*, May 22, 1988; *The Boston Globe*, October 7–8, 2001), in publicly funded universities in Ontario (Anglin and Meng 2000), and—from our personal experience—in universities in Hong Kong.[1] The extent of grade inflation may vary across universities and across departments, but the overall upward trend is unmistakable.

Every college professor probably has his pet theory of grade inflation, but formal models of the phenomenon are few and far between. Most common explanations for grade inflation rely on some kind of "grade illusion" (a cousin of money illusion). For example, one story is that students demand easier grades to help them get into graduate schools or get better jobs, and lazy or untenured professors oblige. This story requires that end-users of grades have static or adaptive expectations, so that they keep being fooled by lowering standards. Moreover, while agency problems no doubt exists within a university, we are not convinced that they provide a complete explanation for grade inflation. Whatever incentives individual instructors or departments may have to give easier grades, they could not have resulted in a systemic inflationary trend without the tacit acceptance, if not support, of the university administration. Occasional admonitions notwithstanding, administrators appear to have done little to reverse the trend, suggesting that the grading policies adopted by instructors or departments are not inconsistent with the objective of the university. Why, then, do universities allow grade inflation to occur?

An interesting perspective on this question is raised in a recent working paper by Ostrovsky and Schwarz (2002). Applying the insight of Crawford and Sobel (1982), they show that if employers use grades as indicators of abilities in setting wages, then, for a wide

[1] For example, in the Faculty of Social Sciences at the University of Hong Kong, the proportion of students who graduated with First Class Honors rose from 1.6 percent to 17.2 percent in the ten-year period between 1990 and 2000.

class of ability distributions, a university can raise the average earnings of its students by adding noise to their transcripts, in effect coarsening the information content of grades. To our knowledge, their paper is the first formal model of grade inflation in which people do not suffer from grade illusion. However, their model explains grade compression more than it does grade inflation. Although grade inflation implies some degree of grade compression (grades cannot be raised beyond $A+$, and the grade distribution becomes concentrated at the upper tail as more and more students get high grades), the reverse is not true. There is no reason that the coarsening of information should take the specific form of grade inflation as opposed to, say, grade deflation.

In this paper we present a simple signaling model of grade inflation. The model has two main ingredients. First, schools observe the abilities of their individual students while employers do not. Second, employers know the distribution of grades in a school, but not the distribution of student abilities within the school. When a school gives a lot of good grades, the labor market cannot fully distinguish whether this is due to an overly liberal grading standard or whether the school is blessed with a large proportion of high-ability students. This gives rise to an incentive for the school to help some of its low-ability students by giving them good grades. However, since employers have rational expectations, this strategy hurts the high-ability students as the value of good grades becomes diluted. A priori, it is not obvious that the terms of such a tradeoff necessarily imply grade inflation. Nevertheless we identify a condition about the objective of the school which ensures that grade inflation is indeed the only equilibrium outcome.

The assumption we make about a school's objective function is that it cares more about its high-ability students than its low-ability students. This assumption serves to provide the "single-crossing condition" in our signaling model. It implies that schools with more high-ability students have greater incentives to give more good grades. Therefore in equilibrium the labor market sees a large percentage of good grades as a signal for a large percentage of high-ability students.

In the "competitive" version of our model, we consider strategic interactions among schools in grading policies that are generated through signaling. To focus on the signaling aspect of the problem, we assume the labor market is sufficiently thick so that schools do not directly compete with one another in placing students in a limited number of job slots. Since wage offers can be conditioned on the individual school of the job applicant and the grading policy of the school, any school that wants to maintain the credibility of its grading policy can do so without worrying about grade inflation in other schools. But if the overall quality of student bodies are correlated (either positively or negatively) across different schools, employers can make inferences about student quality in one school by looking at the general distribution of grades in other schools. We will show that grades and grading policies are "strategic complements" in such a setting, and this creates a channel that makes grade inflation contagious.

It is worth pointing out at that our signaling model applies generally to environments in which raters (e.g., schools) possess private information and care about the rated (e.g., students). End-users of these ratings (e.g., employers) then decipher and make decisions on the basis of information supplied by raters, bearing in mind the latter's information advantage and incentive structure. Consider the case of company audits or stock market recommendations issued by investment banks, for instance. "Chinese walls" notwithstanding, it is not unreasonable to suspect that auditors and investment analysts may have the interest of their clients rather than that of investors at heart. An auditor may have the incentive to declare a risky company healthy, so as to secure other business from this firm. But doing so dilutes the value of the auditor's seal of approval and therefore is not in the interest of its other clients that are financially more sound. As in the case of grade inflation, therefore, it is not obvious that the incentive to relax auditing standards will translate into a lower auditing standard in equilibrium. Our signaling model of grade inflation provides such a link. It turns out, then, that a matter of "academic" interest can shed light on some of the pressing issues in today's business world as well.

2. A Signaling Model of Grade Inflation

A school has two types of students: high-ability (type-H) or "good" students have productivity ω_H; and low-ability (type-L) or "mediocre" students have a lower productivity ω_L. There are two possible states of the world. In state G, the overall quality of the student body is favorable; the proportion ϕ_G of good students is great. In the unfavorable state S, the school has a smaller proportion ϕ_S of good students. The prior probability of state G is π.

A school observes the type of each individual student. This implies that it knows the state of the world (the proportion of good students in the school) as well. The school gives each student either grade A or B. (Yes, C-students are non-existent in this school.) Effectively, a school can send one of two signals: give A to a fraction ϕ_G of students (easy grading, or e), or give A to a fraction ϕ_S of students (tough grading, or t).[2] Note that in the present context easy grading is not synonymous to grade inflation, since there may indeed be a large fraction of good students who deserve an A. Grade inflation is said to occur only when easy grading (e) is chosen in the unfavorable state (S). When grade inflation occurs, all the ϕ_G good students and a measure $\phi_G - \phi_S$ of mediocre students receive grade A. We assume that mediocre students are chosen randomly to receive the inflated grades under grade inflation. Note that we have to distinguish between good students and A-students. Similarly, depending the strategy the school adopts, B-students need not always be mediocre students.

Employers in the labor market observe a student's grade but not his type. Furthermore, they do not know the state of the world. Wage offer to a student can depend on his grade (A or B) and on his school's grading policy (e or t). This amounts to assuming that the proportion of A grades issued by a school is public information. An A grade by itself is meaningless if it is not interpreted in the context of how many As are given.[3] So employers must somehow form expectations about the school's grading policy based on sampling or other methods. To focus on the implications of rational expectations, we adopt the simplest modeling choice by assuming that employers observe e or t. The conclusions of our paper remain unchanged if employers' perception of a school's grading policy contains errors, as long as these errors are unsystematic.[4]

[2] Depending on the level of aggregation desired, a "school" may be interpreted as a professor, a department, or the whole university.

[3] The analogy from price inflation is that nominal prices must be interpreted in the context of the general price level or the money supply.

[4] Our model does not apply when employers have adaptive or other forms of non-rational expectations.

We assume the labor market is competitive in that students are paid their expected productivity. A school maximizes a weighted sum of the wage offers to its type-H and type-L students. This kind of objective function may be justified by altruistic motives or by the need to attract students. Let R be the relative weight on the wage offers to type-H students. A crucial assumption of this model is that $R > 1$; that is, the school cares more about its good students than its mediocre students. Note that if $R = 1$, then *any* grading policy (including random grading) would give the same level of utility to the school since the unweighted sum of wage offers to all students is a constant when employers have rational expectations. If $R < 1$, the interest of the school and the interest of employers would be diametrically opposed, and grades would cease to become useful signals of student ability.[5] Our introspection suggests that the assumption $R > 1$ is probably a reasonable description of the objective function of professors and school administrators. From a more pragmatic point of view, if schools build their reputation and endowments on their more distinguished alumni, it seems natural that they should pay more attention to the wellbeing of their good students than to that of the mediocre ones.

The assumption that $R > 1$ provides the standard "single-crossing" condition in signaling models (Spence 1973). To see this, let $w(A|e)$ and $w(B|e)$ be the wages paid to A and B-students, respectively, when the market observes easy grades given by a school. Similarly, let $w(A|t)$ and $w(B|t)$ be the wages paid to A and B-students when grading is tough. In the favorable state, the difference in payoffs to the school between easy grading and tough grading is

$$
\begin{aligned}
U(e|G) - U(t|G) =& [R\phi_G w(A|e) + (1-\phi_G)w(B|e)] \\
& - [R\phi_S w(A|t) + R(\phi_G - \phi_S)w(B|t) + (1-\phi_G)w(B|t)].
\end{aligned}
$$

In the unfavorable state, the difference in payoffs is

$$
\begin{aligned}
U(e|S) - U(t|S) =& [R\phi_S w(A|e) + (\phi_G - \phi_S)w(A|e) + (1-\phi_G)w(B|e)] \\
& - [R\phi_S w(A|t) + (1-\phi_S)w(B|t)].
\end{aligned}
$$

If employers keep underestimating the extent of grade inflation, a school can reap short term gains for its mediocre students without hurting its good students.

[5] The school would like to give mediocre students A grades and good students B grades. Employers would see through this and assign different meanings to A and B grades. The only equilibrium would be a babbling equilibrium.

The single-crossing condition requires the school to have greater incentive to choose easy grading in the favorable state than in the unfavorable state. In other words, we require

$$[U(e|G) - U(t|G)] - [U(e|S) - U(t|S)] = (R-1)(\phi_G - \phi_S)(w(A|e) - w(B|t)) > 0. \quad (2.1)$$

If, in equilibrium, employers expect that an easy A is better than a tough B (which is indeed the case, as we will demonstrate in the next section), then the assumption $R > 1$ guarantees that the single-crossing condition (2.1) is satisfied.

Intuitively, a school has greater incentive to give more As (choose easy grading) in the favorable state because there are more good students who deserve the grade. Because of this incentive structure, employers rationally use easy grading as a positive signal for the favorable state. This in turn allows the school to engage in some degree of equilibrium grade inflation (choosing easy grading in the unfavorable state), as in a standard signaling model.

Because neither the school nor the market cares about grades per se,[6] our signaling model of grade inflation is what is known as a "cheap talk" game. Strategic information transmission problems arise in our context because the objective of the school and the objective of employers are not completely aligned. The original cheap talk game studied by Crawford and Sobel (1982) focuses on coarsening of informational content in transmission, which in our application would be grade compression as opposed to grade inflation. In our model there are only two states instead of a continuum of them as in Crawford and Sobel, with two possible signals (easy grading or tough grading). These assumptions lead us to concentrate on the issue of grade inflation, although we do not skirt the related issue of compression entirely because inflation in our framework implies compression.

3. Signaling by a Single School

This section deals with signaling by a single school. The analysis applies more generally to the case of more than one schools, provided that the quality of the student body in each school (i.e., the favorable or unfavorable state) is uncorrelated with one another.

[6] The utility function of the market is not explicitly given, but one can imagine that the market minimizes the deviation of the wage offer to a student from his expected productivity.

Because of the incentive structure induced by the single-crossing property, we look for an "inflationary semi-pooling equilibrium," in which the school adopts the following strategy. In the favorable state G, the school chooses e and is "honest" about the quality of its students. It gives all the ϕ_G good students A grades and all the $1-\phi_G$ mediocre students B grades. In the unfavorable state S, the school randomizes between e and t. With probability $1-p$, the school is "honest" and chooses t: only the ϕ_S good students receive A grades. With probability p, the school engages in grade inflation and chooses e. In that case, all ϕ_S good students get A grades, a measure $\phi_G - \phi_S$ of mediocre students also get A grades, and the remaining $1-\phi_G$ mediocre students receive B grades.

Given such a strategy, tough grading is a sure sign of the unfavorable state. The updated probability of state G is zero. Therefore competitive wage offer to A-students is $w(A|t) = \omega_H$ and the wage offer to B-students is $w(B|t) = \omega_L$. When the market observes a large fraction of A grades (easy grading), on the other hand, this could be due to either a large fraction of good students or to grade inflation. Using Bayes' rule, the updated probability of state G is

$$q(G|e) = \frac{\pi}{\pi + (1-\pi)p}. \tag{3.1}$$

Competitive wage offer to A-students is

$$w(A|e) = q(G|e)\omega_H + (1 - q(G|e))\left(\frac{\phi_S}{\phi_G}\omega_H + \left(1 - \frac{\phi_S}{\phi_G}\right)\omega_L\right). \tag{3.2}$$

The wage offer to B-students is $w(B|e) = \omega_L$.

Proposition 3.1. *There is a unique inflationary semi-pooling equilibrium if and only if*

$$\pi < \gamma \equiv \frac{(R-1)\phi_S}{(R-1)\phi_S + \phi_G}, \tag{3.3}$$

with an equilibrium probability of grade inflation equal to

$$p^* = \frac{\pi}{1-\pi}\frac{\phi_G}{\phi_S}\frac{1}{R-1}. \tag{3.4}$$

Proof. The necessary and sufficient conditions for a semi-pooling equilibrium are: (i) in state G the school weakly prefers e to t:

$$R\phi_G w(A|e) + (1-\phi_G)\omega_L \geq R\phi_S\omega_H + R(\phi_G - \phi_S)\omega_L + (1-\phi_G)\omega_L;$$

and (ii) in state S the school is indifferent between e and t:

$$R\phi_S w(A|e) + (\phi_G - \phi_S)w(A|e) + (1-\phi_G)\omega_L = R\phi_S\omega_H + (1-\phi_S)\omega_L.$$

Using equation (3.2) and (3.1), we can solve for the equilibrium probability p^* of grade inflation from the second condition. The solution is given by equation (3.4) as stated in the proposition. Condition (3.3) in the proposition is equivalent to the requirement that $p^* < 1$. Finally, since $R > 1$ and $w(A|e) > w(B|t) = \omega_L$, the single-crossing condition (2.1) discussed in section 2 is satisfied. Therefore, condition (ii) implies condition (i).

Q.E.D.

In the subsequent analysis it is often convenient to characterize the equilibrium directly in terms of the $q(A|e)$ function. To this end, we rewrite the indifference condition between e and t in state S as

$$R\phi_S(\omega_H - w(A|e)) = (\phi_G - \phi_S)(w(A|e) - \omega_L). \tag{3.5}$$

The left-hand-side represents the cost of grade inflation; the right-hand-side represents the benefit. Define

$$\omega_* = \frac{R\phi_S\omega_H + (\phi_G - \phi_S)\omega_L}{R\phi_S + \phi_G - \phi_S} \tag{3.6}$$

as the expected wage for A-students that keeps the school indifferent between e and t in state S. Then the equilibrium indifference condition, equation (3.5), can be simply stated as

$$w(A|e) = \omega_*.$$

Now, define

$$\underline{\omega} = \frac{\phi_S}{\phi_G}\omega_H + \left(1 - \frac{\phi_S}{\phi_G}\right)\omega_L \tag{3.7}$$

as the wage for A-students when the school is believed to have inflated the grades. Then we can express the equilibrium condition as:

$$w(A|e) = q(G|e)\omega_H + (1 - q(G|e))\underline{\omega} = \omega_*,$$

which gives

$$q(A|e) = \frac{\omega_* - \underline{\omega}}{\omega_H - \underline{\omega}}. \tag{3.8}$$

Using equations (3.7) and (3.6), we can show that the expression on the right-hand-side is equal to the term γ defined in the statement of Proposition 3.1. Therefore, the equilibrium indifference condition can be simply stated as:

$$q(G|e) = \gamma. \tag{3.9}$$

Loosely speaking, $q(G|e)$ is the probability in state S of "fooling" the market into believing that the state is favorable. An inflationary semi-pooling equilibrium occurs when the probability p of grade inflation is such that the equilibrium probability $q(A|e)$ of fooling the market is equal to γ. Note that γ is a number between 0 and 1. The function $q(A|e)$ is decreasing in the probability of grade inflation p, with $q(A|e) = 1$ at $p = 0$ and $q(A|e) = \pi$ at $p = 1$. Therefore, a unique solution $p^* \in (0, 1)$ exists for any $\gamma \in (\pi, 1)$.

What happens if $\gamma \leq \pi$? We consider an "inflationary pooling equilibrium," in which the school chooses easy grading in both state G and state S with probability 1.

Proposition 3.2. *There exists an inflationary pooling equilibrium if and only if $\pi \geq \gamma$.*

Proof. We need to prove the school weakly prefers e to t in both state S and state G if $\pi \geq \gamma$. In a pooling equilibrium, since the school always chooses e, the market does not update its probability assessment for state G upon observing easy grading. The competitive wage offer to A-students is

$$w(A|e) = \pi\omega_H + (1 - \pi)\underline{\omega}. \tag{3.10}$$

The wage offer to B-students is ω_L. Let the out-of-equilibrium belief be that the state is S when t is observed. Then the competitive wage offer would be ω_H to A-students and ω_L to B-students. Since $\pi \geq \gamma$, one can verify that the school weakly prefers e to t in state S. Since the single-crossing condition is satisfied, weak preference for e in state S implies strict preference for e in state G.

Consider the reverse statement. In an inflationary pooling equilibrium, when deviation to t occurs, A-students get ω_H while the wage $w(B|t)$ for B-students depends on the out-of-equilibrium belief. However, for any belief, we have $w(B|t) \geq \omega_L$. In equilibrium the school weakly prefers e to t in state S. Therefore,

$$R\phi_S w(A|e) + (\phi_G - \phi_S)w(A|e) + (1 - \phi_G)\omega_L \geq R\phi_S\omega_H + (1 - \phi_S)w(B|t),$$

where $w(A|e)$ is given by equation (3.10). Since $w(B|t) \geq \omega_L$, the above inequality implies that $w(A|e) \geq \omega_*$, from which it follows that $\pi \geq \gamma$. *Q.E.D.*

3.1. Comparative statics

Propositions 3.1 and 3.2 imply that the equilibrium probability of grade inflation (choosing e in state S) is

$$p^* = \min\left\{\frac{\pi}{1-\pi}\frac{\phi_G}{\phi_S}\frac{1}{R-1}, 1\right\}. \tag{3.11}$$

Equation (3.11) shows that a greater concern for the good students lowers grade inflation. Intuitively, a greater R increases the cost of grade inflation (the left-hand-side of equation 3.5) while keeping the benefit unchanged. As a result equilibrium p^* must decrease to increase $w(A|e)$ and restore the indifference condition.

Equation (3.11) also shows that grade inflation increases with the probability of the favorable state. An increase in π affects the cost and the benefit of grade inflation only through its effect on $w(A|e)$. From equation (3.1), a greater π increases the likelihood that easy grading is justified by the favorable state.[7] This tends to raise $w(A|e)$, decreasing the cost and increasing the benefit of grade inflation. To restore the indifference condition, equilibrium grade inflation must occur more often to reduce $w(A|e)$ back to its original level.

An increase in ϕ_G/ϕ_S increases the equilibrium probability of grade inflation. When there is easy grading in the unfavorable state, the fraction of inflated A grades is $1-\phi_S/\phi_G$.

[7] An elite university can claim with a straight face that most of its students get As because they are all good students—often the claim is indeed true. A similar claim made by a lesser school is less convincing.

So the ratio ϕ_G/ϕ_S can be interpreted as the extent of grade inflation: the higher the ratio, the more misleading grades the school reports when grade inflation occurs. In the unfavorable state, the number of mediocre students who benefit from grade inflation is $\phi_G - \phi_S$, and the number of good students who bear the cost of a diminished A grade is ϕ_S. From the indifference condition (3.5), one can see that an increase in ϕ_G/ϕ_S raises the benefit of inflation relative to the cost. Therefore ω_* (the wage for A-students that would keep the school indifferent between easy grading and tough grading in the unfavorable state) falls. To restore the condition $w(A|e) = \omega_*$, equilibrium p^* must increase to reduce $w(A|e)$.[8]

One interesting observation is that grade inflation does not depend on the productivities of high-ability and low-ability students. From equation (3.11), the equilibrium probability p^* is independent of ω_H and ω_L. This is because an increase in productivity premium ω_H/ω_L has two opposing effects. On one hand, from equation (3.6), raising ω_H/ω_L increases ω_*. On the other hand, from equation (3.2), raising ω_H/ω_L increases the $w(A|e)$. These two effects exactly cancel each other.

Does a better school have greater incentives to inflate grades? The answer depends on what "better school" means. If students from a better school are more productive in the sense that ω_H and ω_L are higher, then a better school has the same incentive as other schools to inflate grades. If average student quality (the proportion of good students) is higher in a better school in each state of the world, then whether a better school is more likely to inflate grades depends on the ratio ϕ_G/ϕ_S. An increase in ϕ_G and a decrease in ϕ_S have opposite effects on average student quality, but they both lead to a greater probability of grade inflation.[9] Finally, a "better school" may mean a greater probability

[8] When grade inflation occurs, a larger value of ϕ_G/ϕ_S means that there are more A-students who do not deserve their grade. Hence, other things equal, $w(A|e)$ falls and this will tend to cause the probability of grade inflation to drop to restore equilibrium. Since $R > 1$, the effect of ϕ_G/ϕ_S on $w(A|e)$ is dominated in magnitude by the effect of ϕ_G/ϕ_S on ω_*. Hence the overall effect of an increase in ϕ_G/ϕ_S on the probability of grade inflation is positive.

[9] This result is an artifact of our two-state model with binary student types. In this model, a school can choose to give A grades to either a fraction ϕ_G (easy grading) or a fraction ϕ_S (tough grading) of its students. So changes in ϕ_G and ϕ_S affect not only the quality of the student body, but the meaning of the grading policies (e and t) as well. Other things equal, a higher ϕ_G implies that the school is giving more misleading grades when grade inflation occurs. Given this feature of our model, we prefer to interpret ϕ_G/ϕ_S as a parameter that describes the feasible extent of grade inflation.

of the favorable state G. On average, a school with a larger π has a greater proportion of good students, and our comparative statics result indicates that such a school is likely to inflate its grades more often.

3.2. Welfare analysis

Grade inflation affects the welfare of good and mediocre students differently. In equilibrium he always gets grade A, and receives an expected payoff of

$$\pi w(A|e) + (1-\pi)(p^* w(A|e) + (1-p^*)\omega_H).$$

Since $\omega_H > w(A|e)$, the good student is worse off when there is grade inflation ($p^* > 0$) than he would be if there were no grade inflation. Furthermore, the welfare of the good student falls monotonically with the probability of grade inflation. A higher p^* adversely affects a good student in two ways: it reduces the probability $(1-\pi)(1-p^*)$ of receiving the deserved wage ω_H when the school chooses tough grading, and it reduces the market estimate $w(A|e)$ of the productivity of A-students when the school chooses easy grading.

The opposite is true with a type-L student. He gets grade B in the favorable state and has a probability $(\phi_G - \phi_S)/(1-\phi_S)$ of getting grade A through grade inflation in the unfavorable state. His expected equilibrium payoff is:

$$\pi\omega_L + (1-\pi)\left((1-p^*)\omega_L + p^*\left(\frac{1-\phi_G}{1-\phi_S}\omega_L + \frac{\phi_G-\phi_S}{1-\phi_S}w(A|e)\right)\right).$$

Since $w(A|e) > \omega_L$, a mediocre student is better off when there is grade inflation than what he would otherwise be if there were no grade inflation. Furthermore his welfare rises monotonically with the probability of grade inflation. An increased use of grade inflation improves his chance of receiving a higher wage $w(A|e)$ instead of ω_L, but reduces the wage offer $w(A|e)$ at the same time. The net effect is positive, as the derivative of the equilibrium payoff with respect to p is:

$$(1-\pi)\frac{\phi_G-\phi_S}{1-\phi_S}\left(\underline{\omega} - \omega_L + q(G|e)(\omega_H - \underline{\omega})\frac{\pi}{\pi+(1-\pi)p}\right) > 0.$$

We conclude that the two types of the students have opposing interests with regard to grade inflation.

Although grade inflation hurts good students and benefit mediocre ones, the school is unambiguously worse off in a signaling equilibrium compared to a situation where student ability is public information. With full observability, all good students get A grades and are paid ω_H, while all mediocre students get B grades and are paid ω_L. The school's payoff is then

$$\pi(R\phi_G\omega_H + (1-\phi_G)\omega_L) + (1-\pi)(R\phi_S\omega_H + (1-\phi_S)\omega_L).$$

With imperfect information, the school's expected equilibrium payoff in state G is

$$R\phi_G w(A|e) + (1-\phi_G)\omega_L,$$

and the expected payoff in state S is

$$(1-p^*)(R\phi_S\omega_H + (1-\phi_S)\omega_L) + p^*(R\phi_S w(A|e) + (\phi_G - \phi_S)w(A|e) + (1-\phi_G)\omega_L).$$

In a semi-pooling equilibrium, p^* is such that the school is indifferent between the two signals e and t, so the expected payoff in state S is simply

$$R\phi_S\omega_H + (1-\phi_S)\omega_L.$$

The school is worse off in the semi-pooling equilibrium compared to the case when student ability is public information, because grade inflation lowers the wage offer to its good students in state G from ω_H to $w(A|e)$, without increasing the school's expected payoff in state S. In a pooling equilibrium with $p^* = 1$, the expected payoff in state S is

$$R\phi_S w(A|e) + (\phi_G - \phi_S)w(A|e) + (1-\phi_G)\omega_L,$$

where $w(A|e)$ is given by equation (3.10). A few steps of calculations reveal that the school's equilibrium expected utility is lower than in the case of public information if and only if

$$(R-1)(1-\pi)(1-\phi_S/\phi_G)(\pi\phi_G + (1-\pi)\phi_S)(\omega_H - \omega_L) > 0,$$

which is true because $R > 1$.

The result that the school is worse off in a signaling equilibrium compared to the case when student ability is public information is due to the inability of the school to commit to an honest grading policy. The private incentives to tinker with grades and help a few mediocre students hurt the school because the market understands such incentives and makes adjustments in wage offers accordingly. This result is similar in spirit to the celebrated result about rules versus discretion in the literature on monetary inflation (Kydland and Prescott 1977).

In the present model, employers always pay workers their expected productivities. Their welfare is unaffected by grade inflation. In a more general setting, however, employers care about optimal task assignment according to ability (e.g., Gibbons and Waldman 1999) and about optimal sorting by worker ability (e.g., Kremer 1993). By coarsening the informational content of grades as signals of worker quality, grade inflation will have a negative effect on total output in such kind of environments.

3.3. Equilibrium selection

Because grade inflation hurts good students and benefits mediocre ones, it is conceivable that a school which cares deeply about its good students might want to boost the value of its A grade by limiting its supply (i.e., to engage in grade deflation). Signaling models are often plagued by multiplicity of equilibria, some of which are supported by rather arbitrary out-of-equilibrium beliefs. Though the inflationary (semi-pooling or pooling) equilibrium we have identified survive the standard refinements in game theory, there may be other reasonable equilibrium outcomes in our model. If this is the case, our comparative statics and welfare results would lose much of their force, because it is not clear which of the multiple equilibria would be observed. We therefore devote this subsection to a discussion of the existence of other potential equilibria in our setting.

One of the two possible separating equilibria—truthful grading—can be ruled out immediately. If the school's grading policy truthfully reflects the ability mix of its students, then upon observing easy grading the market would conclude that the state is G, which would prompt the school to inflate grades when the state is actually S.

Next, we rule out the "perverse separating equilibrium" (where the school chooses t in state G and e in state S), the "double-pooling equilibrium" (where it randomizes between e and t in both states), and the "deflationary semi-pooling equilibrium" (where it randomizes between t and e in state G and chooses t in state S). To have any one of these three types of equilibria, the school must (i) weakly prefer e to t in state S; and (ii) weakly prefer t to e in state G. For this to be true, the sign of the single-crossing condition (2.1) has to be reversed. In other words, we require

$$(R-1)(\phi_G-\phi_S)(w(A|e)-w(B|t))\leq 0,$$

which implies that $w(B|t)\geq w(A|e)$. Furthermore, since the school chooses e in state S with positive probability under the proposed equilibria, we have $w(A|e)<\omega_H$. Similarly, since the school chooses t in state G with positive probability under the proposed equilibria, we have $w(B|t)>\omega_L$. These three inequalities imply that

$$R\phi_S w(A|e)+(\phi_G-\phi_S)w(A|e)+(1-\phi_G)\omega_L < R\phi_S\omega_H+(\phi_G-\phi_S)w(B|t)+(1-\phi_G)w(B|t),$$

which contradicts the condition that the school weakly prefers e to t in state S.

Finally, consider a "deflationary pooling equilibrium," in which the school choosing tough grading in both states. In such an equilibrium A-students get $w(A|t)=\omega_H$ and B-students get

$$w(B|t)=\pi\left(\frac{1-\phi_G}{1-\phi_S}\omega_L+\left(1-\frac{1-\phi_G}{1-\phi_L}\right)\omega_H\right)+(1-\pi)\omega_L.$$

If a deviation to e is observed, B-students get ω_L, and the wage $w(A|e)$ for A-students lies between $\underline{\omega}$ and ω_H. This type of pooling equilibria does not exist if in either state the school strictly prefers e to t even when $w(A|e)=\underline{\omega}$. Even when a deflationary pooling equilibrium exists, it is sustained by the market belief that the probability of state G is sufficiently low that the wage $w(A|e)$ paid to A-students when deviation to easy grading is observed is too low to make such deviation profitable. This belief is out of the equilibrium path, and therefore unrestricted by the solution concept of perfect Bayesian equilibrium adopted here. However, we argue that such belief is unreasonable, and the deflationary pooling equilibrium can be ruled out by a standard refinement in signaling games (Banks and Sobel 1987). To show this, it suffices to establish that the lowest wage $w(A|e)$ for A-students that induces a deviation to e is higher in state S than in state G. This condition can be written as:

$$\begin{aligned}&\frac{R\phi_S\omega_H+(\phi_G-\phi_S)w(B|t)+(1-\phi_G)(w(B|t)-\omega_L)}{R\phi_S+\phi_G-\phi_S}\\ >&\frac{R\phi_S\omega_H+R(\phi_G-\phi_S)w(B|t)+(1-\phi_G)(w(B|t)-\omega_L)}{R\phi_S+R(\phi_G-\phi_S)}.\end{aligned} \tag{3.12}$$

One can verify that the above always holds. Equation (3.12) ensures that the set of wage offers that would induce deviation to easy grading under state G strictly contains the set of wage offers that would induce deviation under state S. Thus, the school is "infinitely more likely" to make such deviation in state G than in state S. When deviation to easy grading is observed, the market should then believe that the state is G. Of course, this belief will not support a deflationary pooling equilibrium, as the school would indeed choose to deviate in state G.

Grades being a form of cheap talk in our model, there is always a "babbling equilibrium," in which the school randomly assign grades to students regardless of their ability and the state, and the market ignores grades completely and pays every student the same wage. However, our assumption that $R > 1$ suggests that the school clearly has an incentive to differentiate its good students by giving them A grades. Grades therefore provide a credible language that the school will use to break an agnostic babbling equilibrium. We say that the babbling equilibrium is not *neologism-proof* (Farrell 1993; Farrell and Rabin 1996).

4. Signaling by Two Schools

Our model of signaling by a single school indicates that grade inflation is not merely a problem of "racing to the bottom" in the competition among schools. The phenomenon arises more fundamentally from the inability of schools to commit to an honest grading policy in an environment with private information. Nevertheless introducing strategic interactions among schools in this kind of environment is useful, because it can help us address questions such as: Does competition among schools tend to encourage or constrain grade inflation? How does grade inflation spread from one school to another?

In principle one can consider many kinds of strategic interactions among schools. For example, schools may compete in helping to place their graduates in a fixed number of desirable job slots, or they may compete in trying to attract the most promising incoming students. To focus on the signaling aspect of school competition, we choose to ignore such direct competition. Instead our attention is restricted to an environment in which the labor market is sufficiently thick that all students receive wage offers that are equal to their expected marginal product. In this environment the only channel through which schools interact with one another is the signals they send through grades.

Consider a model in which there are two identical schools, 1 and 2. Each school knows its own state but not the state of the other school. Employers do not observe the state in either school. The correlation structure of the states in the two schools is public information. A simple one-parameter model of correlation is given in the following contingency table:

$$\begin{array}{cc} & \begin{array}{cc} G_1 & \quad S_1 \end{array} \\ \begin{array}{c} G_2 \\ S_2 \end{array} & \left[\begin{array}{cc} \frac{1}{2}\rho & \frac{1}{2}(1-\rho) \\ \frac{1}{2}(1-\rho) & \frac{1}{2}\rho \end{array} \right] \end{array}$$

In this model,

$$\rho = \Pr[G_j | G_i] = \Pr[S_j | S_i],$$

for $i \neq j$ $(i, j = 1, 2)$. The parameter ρ varies between 0 (perfect negative correlation) and 1 (perfect positive correlation). When $\rho = \frac{1}{2}$, the two schools have independent states. Note also that, for simplicity, we have eliminated the parameter π in the single-school model. In the present two-school model, the unconditional probability of the favorable state in either school is:

$$\Pr[G_1] = \Pr[G_2] = \frac{1}{2}.$$

Once again, we consider a symmetric inflationary semi-pooling equilibrium, in which each school $i = 1, 2$ chooses e_i in state G_i and randomizes between e_i and t_i in state S_i. Let p be the common probability that school i chooses easy grading in the unfavorable state.

In such an equilibrium, the conditional probabilities are given by:

$$q(G_1|e_1,e_2) = \frac{\frac{1}{2}\rho + \frac{1}{2}(1-\rho)p}{\frac{1}{2}\rho + (1-\rho)p + \frac{1}{2}\rho p^2};$$
$$q(G_1|e_1,t_2) = \frac{\frac{1}{2}(1-\rho)}{\frac{1}{2}(1-\rho) + \frac{1}{2}\rho p};$$
$$q(G_1|t_1,e_2) = q(G_1|t_1,t_2) = 0.$$

The competitive wage offer to B-students is always equal to ω_L. The wage offer to A-students depends on observed grading policies:

$$w(A_1|e_1,e_2) = q(G_1|e_1,e_2)\omega_H + (1 - q(G_1|e_1,e_2))\underline{\omega};$$
$$w(A_1|e_1,t_2) = q(G_1|e_1,t_2)\omega_H + (1 - q(G_1|e_1,t_2))\underline{\omega};$$
$$w(A_1|t_1,e_2) = w(A_1|t_1,t_2) = \omega_H;$$

where $\underline{\omega}$ is the wage for an A-student when his school is believed to have inflated its grades, as defined by equation (3.7) in the previous section.

We first establish that the single-crossing condition continues to hold in the two-school case.

LEMMA 4.1. *For any $\rho \in (0,1)$, if school $i = 1,2$ weakly prefers easy grading to tough grading in state S_i, then it strictly prefers easy grading to tough grading in state G_i.*

PROOF. Let W_S be school 1's expectation about the wage of its A-students when it chooses easy grading in the unfavorable state. We have

$$W_S = (1-\rho)w(A_1|e_1,e_2) + \rho(pw(A_1|e_1,e_2) + (1-p)w(A_1|e_1,t_2)). \tag{4.1}$$

Similarly, let W_G be school 1's expectation about the wage of its A-students when it chooses easy grading in the favorable state:

$$W_G = \rho w(A_1|e_1,e_2) + (1-\rho)(pw(A_1|e_1,e_2) + (1-p)w(A_1|e_1,t_2)). \tag{4.2}$$

The single-crossing condition requires that $U(e_1|G_1) - U(t_1|G_1) > U(e_1|S_1) - U(t_1|S_1)$. Following the same manipulations as those in Section 2, this can be written as:

$$(R-1)(\phi_G - \phi_S)(W_S - \omega_L) + R\phi_G(W_G - W_S) > 0.$$

Since $R > 1$ and $W_S > \omega_L$, the first term in the above equation is positive. Furthermore one can verify that, for any $\rho \in (0, 1)$,

$$W_G - W_S = (1 - p)(2\rho - 1)(w(A_1|e_1, e_2) - w(A_1|e_1, t_2)) \geq 0.$$

Therefore the single-crossing condition is indeed satisfied. *Q.E.D.*

In the single-school case, the assumption that $R > 1$ helps to ensure that the incentive for a school to choose easy grading is greater in the favorable state than in the unfavorable state. In the present case, there is another reason that the single-crossing condition is satisfied: a school expects its A-students to receive higher expected wages in the favorable state than in the unfavorable state, i.e., $W_G \geq W_S$. With negative correlation, school 1 expects the wage for its A-students to be higher in state G_1 than in state S_1, because in state G_1 school 2 is likely to be in state S_2 and with probability $1-p$ will not inflate. This helps school 1's A-students because when the states are negatively correlated, employers attach a higher probability to G_1 upon observing (e_1, t_2). With positive correlation, the reasoning is the opposite and but reaches the same conclusion. In state G_1, school 2 is likely to be in state G_2 and will for sure have easy grades. This again helps school 1's A-students. Due to positive correlation of the states, employers attach a higher probability to G_1 if the signal is (e_1, e_2).

In an inflationary semi-pooling equilibrium, school 1 must be indifferent between easy grading and tough grading in the unfavorable state. This condition is:

$$(1 - \rho)w(A_1|e_1, e_2) + \rho(pw(A_1|e_1, e_2) + (1 - p)w(A_1|e_1, t_2)) = \omega_*, \tag{4.3}$$

where ω_* is given by equation (3.6) in the previous section. Using the fact that $\gamma = (\omega_* - \underline{\omega})/(\omega_H - \underline{\omega})$, we can write the above equilibrium condition as:

$$f(p) = \gamma, \tag{4.4}$$

where the function $f(p)$ is given by

$$f(p) = (1 - \rho + \rho p)q(G_1|e_1, e_2) + \rho(1 - p)q(G_1|e_1, t_2).$$

Equation (4.4) above reduces to equation (3.9) in the single-school case when $\rho = \frac{1}{2}$ (the prior probability of state G is fixed at $\frac{1}{2}$ in the imperfect correlation model we have constructed.) More generally, for any ρ the equilibrium condition in the two-school case takes the same form as the one in the single-school case. In state S_1, school 1 expects that the other school chooses easy grading with probability $1 - \rho + \rho p$, in which case the market assessment for the probability of state G_1 is $q(G_1|e_1, e_2)$. If the other school chooses tough grading, on the other hand, the market assessment for the probability of state G_1 is $q(G_1|e_1, t_1)$. In any inflationary semi-pooling equilibrium, the probability that school 1 in state S_1 fools the market is equal to γ.

PROPOSITION 4.2. *For any $\rho \in (0, 1)$, there exists an inflationary semi-pooling equilibrium if $\gamma > \frac{1}{2}$, and there exists an inflationary pooling equilibrium if $\gamma \leq \frac{1}{2}$. Furthermore, if $\rho \leq \frac{1}{2}$ and $\gamma > \frac{1}{2}$, the inflationary semi-pooling equilibrium is unique.*

PROOF. For any $\rho \in (0, 1)$, the function $f(p)$ is continuous in p with $f(0) = 1$ and $f(1) = \frac{1}{2}$. Therefore, if $\gamma > \frac{1}{2}$, there exists a $p^* \in (0, 1)$ such that $f(p^*) = \gamma$. Since school $i = 1, 2$ is indifferent between e_i and t_i in state S_i, Lemma 4.1 implies that it will choose e_i in state G_i. Uniqueness of equilibrium is established by the monotonicity of the $f(p)$ function. One can verify that both $q(G_1|e_1, e_2)$ and $q(G_1|e_1, t_2)$ are decreasing in p. Further, $q(G_1|e_1, e_2) \leq q(G_1|e_1, t_2)$ if and only if $\rho < \frac{1}{2}$. Thus, if $\rho \leq \frac{1}{2}$, then $f'(p) < 0$ and the solution to the equation $f(p) = \gamma$ is unique.

In an inflationary pooling equilibrium, each school $i = 1, 2$ chooses e_i in both states with probability 1. Upon observing e_1 and e_2, we have $q(G_1|e_1, e_2) = \frac{1}{2}$. Competitive wage offer to A-students is

$$w(A_1|e_1, e_2) = \frac{1}{2}\omega_H + \frac{1}{2}\underline{\omega}.$$

Let the out-of-equilibrium belief be that school 1's state is S_1 when t_1 is observed. Then the competitive wage offer is ω_H for A-students. One can verify that since $\gamma \leq \frac{1}{2}$, in state S_1 school 1 weakly prefers e_1 to t_1. By Lemma 4.1, weak preference for e_1 in state S_1 implies strict preference for e_1 in state G_1. *Q.E.D.*

Note that negative correlation between states across the two schools is sufficient but not necessary for the uniqueness of equilibrium. For ρ slightly above $\frac{1}{2}$, the function $f(p)$

remains a decreasing function that admits a unique solution to the equilibrium condition $f(p) = \gamma$.

Multiple semi-pooling equilibria can occur when ρ is close to 1. The logic behind the difference between negative and positive correlation is entirely intuitive in our setup. In both cases, the condition for an inflationary semi-pooling equilibrium is that the probability p of grade inflation is such that the probability $f(p)$ of fooling the market is equal to the exogenous parameter γ. Under negative correlation, $f(p)$ decreases with p both because the market becomes more skeptical of good grades ($q(G_1|e_1, e_2)$ and $q(G_1|e_1, t_2)$ decrease) and because it is more likely that the market observes easy grades in both schools, which is bad news because the states are negative correlated ($q(G_1|e_1, e_2) < q(G_1|e_1, t_2)$ when $\rho < \frac{1}{2}$). Monotonicity of $f(p)$ means that there can be at most one semi-pooling equilibrium. The situation is different with positive correlation. A greater p still means a more skeptical market, but having easy grades in both schools is now good for the schools. As a result $f(p)$ can increase over some range of p. Indeed, when ρ is close to 1, the function $f(p)$ is non-monotone so that multiple semi pooling equilibria occur for γ just below and/or just below $\frac{1}{2}$. See Figure 1.

4.1. Grade inflation and correlation

In this subsection, we examine how the degree of correlation of states across the two schools affects the equilibrium probability of grade inflation. Positive correlation in the quality of the student bodies across schools may arise because general economic conditions (e.g., business cycles, the size of the skills premium) affect the decision to enter college, or because the overall environment of teaching and research affects the value-added of the education process. On the other hand, competition by schools for the same cohort of good students may result in a negative correlation of states. Is grade inflation more or less likely when states become more (positively or negatively) correlated?

We restrict our attention to the case $\gamma > \frac{1}{2}$, so that an inflationary semi-pooling equilibrium always exists. The condition for the equilibrium probability p^* of grade inflation is $f(p^*; \rho) = \gamma$. If p^* is unique, then $\partial p^*/\partial\rho$ has the same sign as $\partial f/\partial\rho$. If there are multiple solutions to the equilibrium condition, then the same conclusion applies when p^*

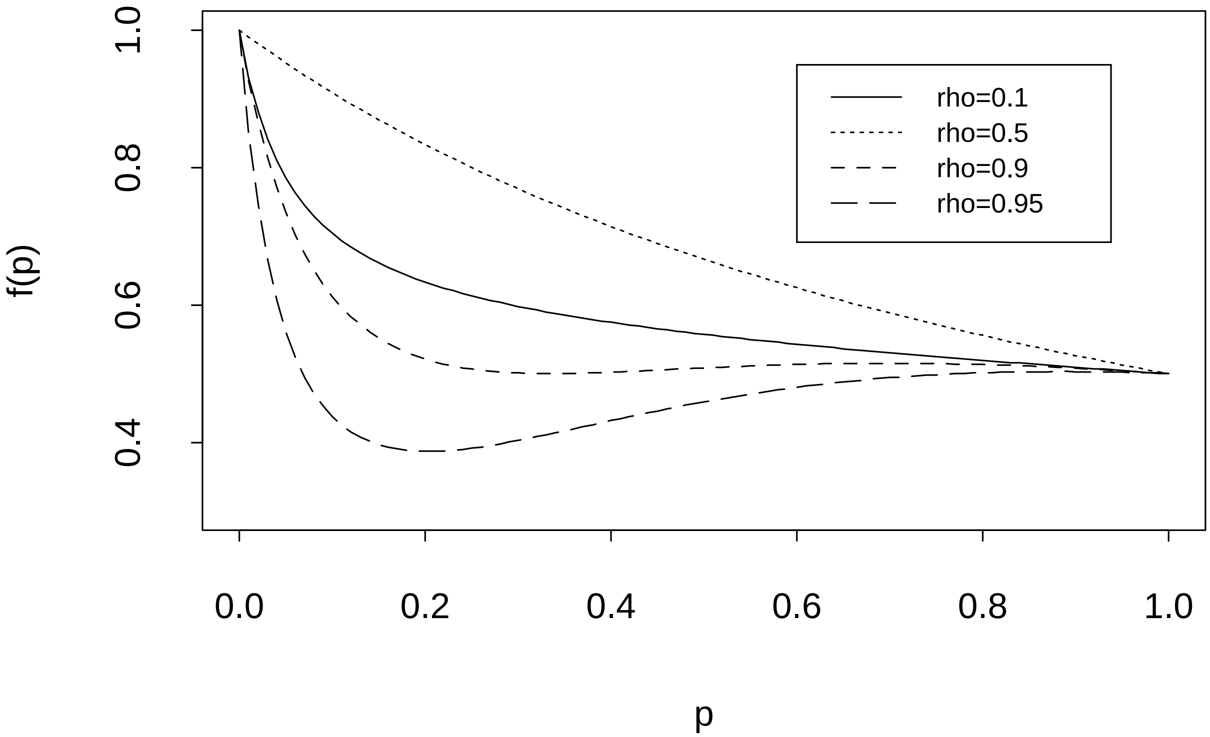

Figure 1

is the largest or the smallest equilibrium solution (Milgrom and Roberts 1994). A direct calculation yields:

$$\frac{\partial f(p;\rho)}{\partial \rho} = \frac{(1+(3-2\rho)p+2\rho p^2)(1-p)p(1-2\rho)}{(\rho+2(1-\rho)+\rho p^2)^2(1-\rho+\rho p)^2}.$$

Thus, $\partial f(p;\rho)/\partial\rho$ has the same sign as $1-2\rho$. If $\rho \leq \frac{1}{2}$, a decrease in ρ (increased negative correlation) shifts the $f(p)$ curve down and reduce the equilibrium probability of grade inflation. If $\rho \geq \frac{1}{2}$, an increase in ρ (increased positive correlation) also shifts the $f(p)$ curve down and reduce grade inflation. See Figure 1. Therefore, the equilibrium probability of grade inflation reaches a maximum at $\rho = \frac{1}{2}$, when the two schools have independent states. Greater correlation in states across the two schools, either negative or positive, serves to constrain grade inflation.

Correlation constrains grade inflation because it reduces the probability that each school fools the market with easy grades. A more positive correlation has two opposing effects: it improves the credibility of bilateral easy grades (increases $q(G_1|e_1,e_2)$ for school

1) and reduces the credibility of unilateral easy grades (decreases $q(G_1|e_1, t_2)$ for school 1). However, a deciding third effect of an increase in positive correlation is that it makes it more likely to for employers to observe t_2 under state S_1. This makes fooling the market more difficult because under positive correlation unilateral easy grading is not as credible as bilateral easy grading ($q(G_1|e_1, t_2) < q(G_1|e_1, e_2)$ when $\rho > \frac{1}{2}$). A more negative correlation also makes it more difficult to fool the market, but for the opposite reason: it makes less likely for employers to observe t_2 under S_1, and under negative correlation unilateral easy grades is the more credible signal.

The case of independent states ($\rho = \frac{1}{2}$) is equivalent to the single-school case analyzed in section 3. It might appear against one's intuition that grade inflation is generally less serious with two schools than with a single school. But in our framework schools do not directly compete with one another in placements or in any other way. The only strategic interaction between schools comes from the underlying inference problem faced by the labor market when the schools have correlated states. Grade inflation is less serious with two schools because it is harder to fool the market when there are two signals available instead of one.

4.2. Strategic interactions

We know from the single-school case that grade inflation at an isolated school can result from a number of exogenous changes, such as a smaller R, a greater π, or a greater ϕ_G/ϕ_S. How do changes in one school's grading policy affect other schools? How does the new equilibrium compare with the one before the changes took place? The answers depend on whether grade inflation policies at different schools are strategic complements or strategic substitutes.

For each school $i = 1, 2$, denote the equilibrium probability of inflation at school i by p_i^*. The indifference condition for school 1 is

$$f_1(p_1^*, p_2^*) = \gamma_1, \tag{4.5}$$

where γ_1 is defined as in equation (3.3) using the parameter values for school 1. The function f_1 is given by

$$f_1(p_1, p_2) = (1 - \rho + \rho p_2)q(G_1|e_1, e_2) + \rho(1 - p_2)q(G_1|e_1, t_2),$$

where

$$q(G_1|e_1, e_2) = \frac{\rho + (1-\rho)p_2}{\rho + (1-\rho)(p_1+p_2) + \rho p_1 p_2},$$

and

$$q(G_1|e_1, t_2) = \frac{1-\rho}{1-\rho+\rho p_1}.$$

We will think of equation (4.5) as defining the "reaction function" $p_1^*(p_2)$ for school 1. For any $\gamma_1 > \frac{1}{2}$, we have $f_1(0,0) = 1 > \gamma_1$, implying $p_1^*(0) > 0$. Similarly, $f_1(1,1) = \frac{1}{2} < \gamma_1$, which implies $p_1^*(1) < 1$. The same conclusion holds for school 2. Therefore, for any $\gamma_1, \gamma_2 \in (\frac{1}{2}, 1)$, there exists one least one inflationary semi-pooling equilibrium.

The slope of the reaction function $p_1^*(p_2)$ is

$$\frac{dp_1^*}{dp_2} = -\frac{\partial f_1/\partial p_2}{\partial f_1/\partial p_1}.$$

For any ρ, we have:

$$\frac{\partial f_1}{\partial p_1} = (1-\rho+\rho p_2)\frac{\partial q(G_1|e_1,e_2)}{\partial p_1} + \rho(1-p_2)\frac{\partial q(G_1|e_1,t_2)}{\partial p_1} < 0,$$

since both $q(G_1|e_1, e_2)$ and $q(G_1|e_1, t_2)$ are decreasing functions of p_1. Also,

$$\begin{aligned}\frac{\partial f_1}{\partial p_2} &= \rho(q(G_1|e_1,e_2) - q(G_1|e_1,t_2)) + (1-\rho+\rho p_2)\frac{\partial q(G_1|e_1,e_2)}{\partial p_2} \\ &= \frac{p_1(1-2\rho)^2}{(\rho+(1-\rho)(p_1+p_2)+\rho p_1 p_2)^2(1-\rho+\rho p_1)} \geq 0.\end{aligned}$$

It follows that $dp_1^*/dp_2 \geq 0$, and p_1 and p_2 are strategic complements.[10]

An increased use of grade inflation by school 2 makes the event of both schools having easy grades more likely. Under positive correlation, this helps school 1 fool the market, because having easy grading at both schools is a stronger signal for state G_1 than having easy grading at school 1 only ($q(G_1|e_1, e_2) > q(G_1|e_1, t_2)$ when $\rho > \frac{1}{2}$). But the second effect is that it changes the market estimate of state G_1 when both schools have easy grades. Under positive correlation, an increase in p_2 lowers $q(G_1|e_1, e_2)$ because the market will put more weight on the event that both schools are inflating grades. This makes it more

[10] If the states in the two schools are uncorrelated ($\rho = \frac{1}{2}$), then p_1 and p_2 are strategically independent.

difficult for school 1 to fool the market. The first effect dominates the second effect under positive correlation, and the net effect of an increase in p_2 is an increase in $f_1(p_1, p_2)$. Under negative correlation, the sign as well as the relative magnitude of these two effects are reversed, so the net result is still $\partial f_1/\partial p_2 > 0$.

Therefore, under both positive and negative correlation, it is easier for school 1 to fool the market with grade inflation with an increased use of grade inflation by school 2. The result is that grade inflation is a supermodular game between the two schools. Suppose there is a decrease in γ_2 (resulting from, say, a decrease in school 2's concern R_2 for its good students) while γ_1 remains unchanged, school 2 will raise its probability of grade inflation in response to this change. As p_2 rises, even though there is no change in the underlying parameters at school 1, school 1 finds that it is easier to fool the market into believing that its state is favorable when it gives easy grades. In equilibrium, both p_1^* and p_2^* will rise.[11] Thus, strategic interactions between the schools provide a mechanism through which grade inflation is transmitted from one school to another.

5. Signaling by Many Schools

The probability model we use in Section 4 is suitable for studying how the degree of correlation in student quality across two schools affects grade inflation; it is less flexible for studying the case of more than two schools. Extending the analysis to signaling by many schools is useful because it can provide comparative statics and limit results for changes in the number of schools in the system.

We build a stylized model with common shocks and idiosyncratic shocks.[12] Let there be two aggregate states, G and S, with prior probabilities θ and $1-\theta$, respectively. In state S, each school i ($i = 1, \ldots, N$) is in state S_i with probability 1. They all have a student body with a small proportion ϕ_S of good students. In state G, each school i has

[11] If there are multiple equilibria, the same conclusion applies to the maximum and the minimum equilibria (Milgrom and Roberts 1990).

[12] This type of model necessarily implies positive correlation across any two schools. It is generally more difficult to construct models with "negative correlation" when there are three or more schools because negative correlation between any pair of schools is not a transitive relationship.

an independent probability π of reaching state G_i (with a greater proportion ϕ_G of good students) and a probability $1-\pi$ of reaching state S_i. Each school knows its state, but not those of other schools. This model includes the single-school case analyzed before as a special case, with $N=1$ and $\theta=1$. If we let $N=2$, $\pi=\rho$, $\theta=1/(2\rho)$, then it becomes the two-school model with positive correlation.

Consider an inflationary semi-pooling equilibrium, in which each school i chooses easy grading e_i with probability 1 in state G_i and with probability p in state S_i. As in Sections 3 and 4, equilibrium is characterized by the condition $f(p)=\gamma$, where γ is a function of the exogenous parameters as defined in equation (3.3) above, and the function $f(p)$ is school i's assessment of how easy it is for it to "fool" the market into believing that the state is G_i when e_i is observed. We proceed to derive an explicit expression for $f(p)$.

If the aggregate state is favorable (G), then the total probability that any one school will choose easy grading is $\pi+(1-\pi)p$. If the aggregate state is unfavorable (S), then the probability that any one school will choose easy grading is just p. Conditional on state S_i, the probability of the aggregate state being G is given by

$$\mu=\frac{\theta(1-\pi)}{\theta(1-\pi)+1-\theta}.$$

Therefore, if school i chooses easy grading in state S_i, its assessment of the probability that the market observes a total of k ($k=1,\ldots,N$) schools with easy grades is

$$\Pr[k|S_i]=\mu b(N-1,k-1,\pi+(1-\pi)p)+(1-\mu)b(N-1,k-1,p),$$

where $b(\cdot,\cdot,\cdot)$ denotes the binomial probability function (i.e., $b(N,k,p)$ is the probability of observing k successes out of N Bernoulli trials with independent probability of success p). Let $q(G_i|k)$ be the market's assessment of the probability of state G_i when there are k schools (including school i) that chooses easy grading. Then, applying Bayes' rule, we have

$$q(G_i|k)=\frac{\theta b(N,k,\pi+(1-\pi)p)}{\theta b(N,k,\pi+(1-\pi)p)+(1-\theta)b(N,k,p)}\frac{\pi}{\pi+(1-\pi)p}.$$

The first fraction on the right-hand-side is the probability that the aggregate state is G given k schools with easy grades; the second fraction is the probability that school i is in the favorable state given that the aggregate state is favorable. The total probability of fooling the market into believing the state is G_i is given by

$$f(p) = \sum_{k=1}^{N} \Pr[k|S_i] q(G_i|k).$$

The function $f(p)$ is continuous with $f(0) = 1$ and $f(1) = \theta\pi$. Therefore for any finite N, an inflationary semi-pooling equilibrium exists if $\gamma > \theta\pi$. One can also show that an inflationary pooling equilibrium exists if $\gamma \leq \theta\pi$.

Analytical results for the effect of changing N are difficult to obtain because the $f(p)$ function involves combinatorials. We have tried a variety of parameter values for θ, π and N and are able to verify numerically that, for all p,

$$f(p; N+1) < f(p; N).$$

Figure 2 illustrates (using the parameter values $\theta = 0.6$ and $\pi = 0.4$). The presence of more schools, and hence more independent signals, makes it more difficult for any individual school to "fool" the market. Consequently, a larger N causes the equilibrium probability p^* of grade inflation to fall.[13]

If an increase in N reduces grade inflation, will grade inflation be eliminated when there is an arbitrarily large number of schools? In other words, will the equilibrium probability of grade inflation converge to 0, or will it converge to a limit bounded away from 0? When N is arbitrarily large, in a semi-pooling equilibrium, the proportion of schools with easy grades is p in the aggregate state S, and is $\pi + (1 - \pi)p$ in the aggregate state G. Thus, the market can perfectly infer the aggregate state from the proportion of schools with easy grades. When school i observes that its own state is unfavorable S_i, it infers that there is a probability μ that the aggregate state is favorable. In that case, the market observes a fraction $\pi + (1 - \pi)p$ of the schools with easy grades, knowing that only a fraction π of the schools are truly in the favorable state. Therefore, the market assigns a probability

$$\Pr[G_i|G] = \frac{\pi}{\pi + (1-\pi)p}$$

[13] Note that if there are multiple equilibria, this conclusion applies to the largest and the smallest equilibrium p^*.

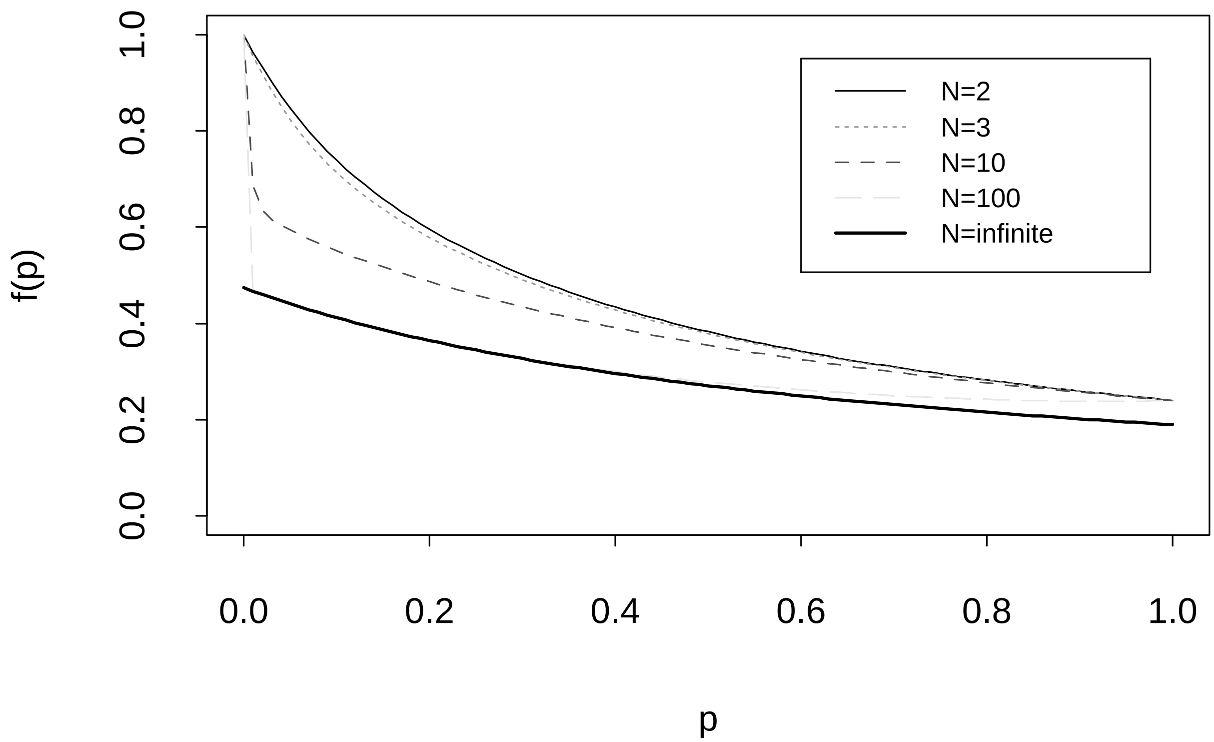

Figure 2

that the student mix in school i is favorable. The school also perceives that with probability $1-\mu$, the aggregate state is unfavorable. In that case, the market can tell with certainty that the aggregate state is unfavorable, and assigns a probability $\Pr[G_i|S] = 0$ that the student mix in school i is favorable. Thus, the probability that school i can fool the market into believing that its student mix is favorable by inflating its grades is

$$f_*(p) = \mu \Pr[G_i|G] + (1-\mu)\Pr[G_i|S] = \frac{\mu\pi}{\pi + (1-\pi)p}.$$

For N arbitrarily large, the condition for an inflationary semi-pooling equilibrium is given by the equation $f_*(p) = \gamma$. The function $f_*(p)$ is decreasing in p, with $f_*(0) = \mu$ and $f_*(1) = \mu\pi$. The same arguments as in Proposition 3.1 and Proposition 4.2 can be used to establish the following:

Proposition 5.1. *For N arbitrarily large, an inflationary semi-pooling equilibrium exists if $\gamma \in (\mu\pi, \mu)$.*

What happens when $\gamma \geq \mu$? In that case, the probability of fooling the market for school i in state S_i is lower than what it takes for it to be indifferent between e_i and t_i, even if in equilibrium no school inflates grades. We then have a separating equilibrium, with honest grading by each school.[14] Therefore, an increase in number of schools can potentially eliminate grade inflation when γ is large enough (say, because R is large). However, if γ is less than μ, then equilibrium grade inflation will persist even in the limit as the number of schools grows indefinitely.

If γ is relatively small, then, in an inflationary pooling equilibrium. the observed proportion of schools with easy grades is 1 in both state G and state S, so the market cannot distinguish the two aggregate states. The competitive wage offer for A-students from any school with easy grades is then

$$w(A|G,e) = \theta\pi\omega_H + (1-\theta\pi)\underline{\omega}.$$

The probability of fooling the market with grade inflation is $\theta\pi$, instead of $f(1) = \mu\pi$. This discontinuity in the probability of fooling the market at $p = 1$ is due to the fact that the market is able to distinguish the two aggregate states except when $p = 1$. The result is that there exists an inflationary pooling equilibrium if and only if $\gamma \leq \theta\pi$. Note that since $\mu < \theta$, if $\mu\pi < \gamma < \theta\pi$, there exist both a semi-pooling equilibrium and a pooling equilibrium.

6. Curbing Grade Inflation

If grade inflation garbles the signaling value of grades and reduces the welfare of the school, are there feasible ways to curb it? In this section, we use our model to discuss two methods that have been proposed (and adopted by some schools) to tame grade inflation.

Some universities have experimented with putting two grades on student transcripts: the student's individual grade and the class average grade (*The Economist*, April 12, 2001). The idea is to provide employers with more information to evaluate the meaning of any

[14] For finite N, an inflationary semi-pooling equilibrium exists even in this case. However, the equilibrium probability of grade inflation is arbitrarily close to zero when N is large.

individual job applicant's grades. If employers have static or adaptive expectations about a school's grading policy, such a move can reduce their "grade illusion" and hence remove the school's temptation to inflate grades. Although we do not dispute the value of providing more information on transcripts, we doubt if this can solve the problem of grade inflation. Grade inflation occurs because sometimes easy grades *are* justified. A high class average grade does not immediately imply lax grading; perhaps the school just has a lot of good students. In our model, employers can perfectly observe the class average grade (i.e., the percentage of As and Bs), yet equilibrium grade inflation persists.

Another strategy to tame grade inflation is to assess students strictly on the basis of ranks. Some universities, for example, fix the proportion of their students graduating with honors (*The New York Times*, May 22, 1988), so that a good student cannot receive honors if his peers are outstanding. This kind of policy requires commitment, since we have shown in Section 3.3 that tough grading is not an equilibrium in our signaling model. Nevertheless, committing to a fixed proportion of As is a lot easier than committing to honest grading. The latter requires varying the proportion of As with the underlying state, which is unverifiable. The former only requires the school to give the same proportion of As every year. The school's reputation can suffer if it breaks its commitment.

Committing to a fixed proportion of As is not a first-best policy, since the proportion of good students changes from year to year. It is therefore interesting to see whether the commitment policy is better than the equilibrium policy with grade inflation from the school's perspective. Let $U(\phi)$ be the school's payoff when it commits to giving a fixed proportion ϕ of A grades. Naturally, ϕ lies between ϕ_S and ϕ_G. Under this policy, the competitive wage for A-students is

$$w(A) = \pi\omega_H + (1-\pi)\left(\frac{\phi_S}{\phi}\omega_H + \left(1 - \frac{\phi_S}{\phi}\right)\omega_L\right).$$

The wage for B-students is

$$w(B) = \pi\left(\left(1 - \frac{1-\phi_G}{1-\phi}\right)\omega_H + \frac{1-\phi_G}{1-\phi}\omega_L\right) + (1-\pi)\omega_L.$$

Then,

$$\begin{aligned} U(\phi) =& \pi(R\phi w(A) + R(\phi_G - \phi)w(B) + (1-\phi_G)w(B)) \\ &+ (1-\pi)(R\phi_S w(A) + R(\phi - \phi_S)w(B) + (1-\phi)w(B)). \end{aligned}$$

Observe that as ϕ increases, both $w(A)$ and $w(B)$ decrease but the coefficient of $w(A)$ in $U(\phi)$ increases, so there are two opposing effects. A few steps of calculations show that $U'(\phi)$ has the same sign as

$$\frac{\pi}{1-\pi} - \frac{\phi_S}{\phi}\frac{1-\phi}{1-\phi_G}.$$

If $U(\phi)$ has a stationary point in the range $[\phi_S, \phi_G]$, the stationary point is a local minimum. It follows that the optimal commitment level ϕ is either ϕ_S or ϕ_G.

Would the school be better off with a commitment to tough grading or easy grading? To answer the question, we first note that when the school is in an inflationary pooling equilibrium, its payoff is just $U(\phi_G)$. Also, the payoff in an inflationary semi-pooling equilibrium is greater than $U(\phi_G)$, because in such equilibrium the school is indifferent between easy grading and tough grading in state S, with an equilibrium wage to A-students higher than that with a commitment to ϕ_G. Thus, we only need to compare the equilibrium payoff U^* to $U(\phi_S)$. A few steps of calculations reveal that $U^* < U(\phi_S)$ if and only if

$$(1-\pi)(1-q(A|e))\frac{\phi_S}{\phi_G} - \pi^2\frac{1-\phi_G}{1-\phi_S} > (q(A|e)-\pi)\left(\pi + \frac{1}{R-1}\right).$$

In an inflationary pooling equilibrium $q(A|e) = \pi$ so the above condition becomes

$$(1-\pi)^2\frac{\phi_S}{\phi_G} > \pi^2\frac{1-\phi_G}{1-\phi_S}. \tag{6.1}$$

In an inflationary semi-pooling equilibrium, using the definition of $q(A|e)$ (equation 3.1) we have instead:

$$\frac{\phi_G}{(R-1)\phi_S + \phi_G} > \left(1 - \frac{1}{R}\right)\left(\pi\frac{1-\phi_G}{1-\phi_S} + 1 - \pi\right). \tag{6.2}$$

In either type of equilibrium, it is possible that commitment to tough grading makes the school strictly better off.[15] However, the conditions required for the two types of equilibria tend to go the opposite directions of (6.1) and (6.2). In a pooling equilibrium,

[15] Since equation (6.1) is independent of R, given any π that satisfies it, by equation (3.4) we can always choose R sufficiently close to 1 to have a pooling equilibrium. Further, observe that for any fixed π equation (6.2) reduces to (6.1) if R is such that $p^* = 1$ by equation (3.4). Thus, for any π that satisfies equation (6.1), we can have a semi-pooling equilibrium that satisfies equation (6.2) by choosing any R that is just greater than the one that makes $p^* = 1$ by equation (3.4).

grade inflation is at the highest level. If this results from the high prior probability of the favorable state, then commitment to tough grading can hurt the school: an increase in π tends to reverse (6.1). Intuitively, commitment to tough grading can be too costly because it often forces the school to give Bs to some of its good students. In a semi-pooling equilibrium, grade inflation is limited due to a lower likelihood of the favorable state or a greater concern for good students. The same two factors tend to reverse the inequality of (6.2). In this case the benefit to the school from the commitment to tough grading is small relative to its cost, and the school will not make such commitment even if it is credible.

Given that schools may not have the incentive to self-discipline, external inducements may be necessary to curb grade inflation, which would be worthwhile if, for example, sorting of workers is important for economic efficiency. One way of discouraging grade inflation is to align the interest of the school more closely with the welfare of its good students (increase R in our model). This may be achieved by tying a school's funding more closely to the long-term labor market performance of its students. Greater tax benefits for alumni donations or increased public funding for schools for outstanding achievements of their alumni will make it more in the interest of the school to ensure that the abilities of the students are duly recognized by the market. Over the long run, this can encourage more honest grading in schools.

7. Summary and Discussions

If grades convey information about the *relative* merits of students, why would schools have the incentive to adopt overly liberal grading standards? The answer we propose is that employers cannot fully distinguish between a situation in which a school is giving lots of easy As and a situation in which the school simply has many good students. Indeed, because a school with more good students has an incentive to give more As, employers use a liberal grading curve as a signal to infer high overall student quality in the school. This does not imply that grades will shoot to the roof, since schools also care about preserving the value of As for its good students. We identify an equilibrium level of grade inflation in this paper, and are able to show that honest grading or other kinds of grading strategies (such as random grading or grade deflation) are not reasonable equilibrium outcomes in our setting.

Grade inflation helps mediocre students at the expense of good students. In a manner similar to the central bank dilemma (Kydland and Prescott 1977), schools would gain if they could commit to an honest grading policy. Our comparative statics results show that schools with a smaller concern for its good students or schools with a higher chance of having a large fraction of good student tend to engage in grade inflation more often.

In an environment with more than one schools, we show that grade inflation by one school makes it easier for another school to fool the market with grade inflation. Thus grade inflation policies are strategic complements, and this provides a channel that makes grade inflation contagious. Nevertheless the availability of signals from other schools does reduce the equilibrium level of grade inflation in our setup.

Grade inflation is a multi-faceted phenomenon, and we do not pretend to have covered all grounds in this paper. Our model is an equilibrium model with no dynamics. It explains why grades are too high, but says little about why they keep rising. One possible approach to model the dynamics of grade inflation is to introduce ambiguity and reputation effects as in central bank models of monetary inflation (Cukierman and Meltzer 1986; Rogoff 1987). We believe that the tradeoff between helping mediocre students and hurting good students, as well as the signaling constraints needed to sustain equilibrium, will provide the basic building blocks of a dynamic model.

To focus on the implications of rational expectations, one abstraction we have impose is that employers have perfect information about the grading policy of a school (that is, they observe e or t without error). We show that even in this case, there is room for equilibrium grade inflation to occur. Nevertheless, introducing imperfect information about a school's grading policy may bring another incentive to inflate grades. As in Lucas's (1973) model of aggregate supply, employers may not be able to distinguish between local shocks (grades on the few transcripts they see) from aggregate shocks (the schools grading curve). Then, schools may be able to reap short term gains from practicing grade inflation.

In this paper we focus our attention to the signaling aspect of competition. But schools compete in other dimensions too. For example, if there are rents in a fixed number of desirable positions, good students may benefit disproportionately from a tight grading policy in the competition for these positions. Furthermore, since grading policies affect the relative well-being of different types of students, they have implications for the competition for incoming students and for the sorting of students by schools as well. These interesting questions have not been addressed in our present work.

Finally, we have not looked into issues related to the role of grades as motivator. A paper by Costrell (1994) studies how educational standards affect students' incentives to exert effort. In that paper, lower standards (grade inflation) can reduce the effort of good students while raising the effort of the marginal students. Incorporating student effort into our model probably reinforces our conclusion about the welfare effects of grade inflation.

References

Anglin, Paul M.; and Meng, Ronald. "Evidence on Grades and Grade Inflation at Ontario's Universities." *Canadian Public Policy* 26 (September 2000): 361–368.

Banks, J.; and Sobel, J. "Equilibrium Selection in Signaling Games." *Econometrica* 55 (1987): 647–662.

The Boston Globe. "Matters of Honor." October 7–8, 2001.

Carney, P.; Isakson, R.L.; and Ellsworth. "An Exploration of Grade Inflation and Some Related Factors in Higher Education." *College and University* 53 (1978): 217–230.

Costrell, Robert M. "A Simple Model of Educational Standards." *American Economic Review* 84 (September 1994): 956–971.

Crawford, Vincent; and Sobel, Joel. "Strategic Information Transmission." *Econometrica* 50 (1982): 1431–1451.

Cukierman, Alex; and Meltzer, Allan H. "A Theory of Ambiguity, Credibility, and Inflation under Discretion and Asymmetric Information." *Econometrica* 54 (November 1986): 1099–1128.

The Economist. "All Shall Have Prizes." April 12, 2001.

Farrell, Joseph. "Meaning and Credibility in Cheap-Talk Games." *Games and Economic Behavior* 5 (1993): 514–531.

Farrell, Joseph; and Rabin, Matthew. "Cheap Talk." *Journal of Economic Perspectives* 10 (Summer 1996): 103–118.

Gibbons, Robert; and Waldman, Michael. "A Theory of Wage and Promotion Dynamics Inside Firms." *Quarterly Journal of Economics* 114 (November 1999): 1321–1358.

Juola, Arvo E. "Illustrative Problems in College-Level Grading." *Personnel and Guidance Journal* (September 1968): 29–33.

Kolevzon, M.S. "Grade Inflation in Higher Education: A Comparative Study." *Research in Higher Education* 15 (1981): 195–212.

Kremer, Michael. "The O-Ring Theory of Economic Development." *Quarterly Journal of Economics* 108 (August 1993): 551–575.

Kydland, Finn; and Prescott, Edward. "Rules Rather Than Discretion: The Inconsistency of Optimal Plans." *Journal of Political Economy* 85 (June 1977): 471–491.

Lucas, Robert E. "Some International Evidence on Output-Inflation Tradeoffs." *American Economic Review* 63 (June 1973): 326–334.

McKenzie, Richard B. *The Political Economy of the Education Process*. Boston: Martinus Nijhoff, 1979.

Milgrom, Paul; and Roberts, John. "Rationalizability, Learning and Equilibrium in Games with Strategic Complementarities." *Econometrica* 58 (November 1990): 1255–1278.

Milgrom, Paul; and Roberts, John. "Comparing Equilibria." *American Economic Review* 84 (June 1994): 441–459.

Millman, J.; Slovaced, S.P.; Kulick, E.; and Mitchell, K.J. "Does Grade Inflation Affect the Reliability of Grades?" *Research in Higher Education* 19 (1983): 423–429.

The New York Times. "Yale Moves to Make *Cum Laude* Mean More." May 22, 1988.

Ostrovsky, Michael; and Schwarz, Michael. "Equilibrium Information Disclosure: Grade Inflation and Unraveling." Working paper, Harvard University, 2002.

Rogoff, Kenneth. "Reputational Constraints on Monetary Policy." *Carnegie-Rochester Conference Series on Public Policy* 26 (1987): 141–181.

Sabot, R.; and Wakermann-Linn, J. "Grade Inflation and Course Choice." *Journal of Economic Perspectives* 5 (Winter 1991): 159–170.

Spence, Michael. "Job Market Signaling." *Quarterly Journal of Economics* 87 (1973): 355–374.