An Economic Theory of Grade Inflation

by Huanxing Yang and Chun Seng Yip [*]

First Draft: April 26, 2002
This Version: November 4, 2003

Abstract

We present a formal analysis of grade inflation. In a labor market where firms must rely on job applicants' college transcripts to assign them to jobs, universities can choose to grade-inflate, i.e., give good grades to its bad students, thus helping them secure better jobs. In doing so the university ignores its impact on the average quality of students with good grades. Grade inflation thus arises due to a free-rider problem, and is contagious: thus "bad" grades drive out "good." Furthermore, grade inflation is exacerbated by (1) university reputation and (2) skills differentials. Finally welfare analysis suggests that universities, firms and students are worse off in a world with grade inflation.

Keywords: Grade Inflation, Higher Education, Adverse Selection, Free-Riding, Collective Reputation. **JEL:** D82, I21, J71.

[*]Department of Economics, University of Pennsylvania, PA 19104. We wish to thank Ken Burdett, Sam-Ho Lee, Nicola Persico and Rafael Rob for helpful conversations and suggestions. All errors are ours.

I Introduction

In recent years, the phenomenon of grade inflation in higher education [1] has been a subject of widespread public interest. Although not a new issue, recent interest in the topic was aroused by the allegation that some of the nation's more prestigious universities were most susceptible to inflated transcripts compared to other institutions. The Boston Globe reported in October 2001 that a staggering 91 percent of Harvard University's class of 2001 graduated with honors (Boston Globe, 2001a, b). The newspaper also revealed that high rates of honors were prevalent across many Ivy League universities and not just at Harvard. Table 1, taken from the Globe, shows the proportion of seniors graduating with honors at these colleges in 2001, while contrasting them with some non-Ivy League universities.

(table 1 about here)

The Ivy League is not alone. Sabot and Wakeman-Linn (1991) provide evidence of rising grades among a sample of eight prominent undergraduate programs including highly ranked liberal arts colleges. [2] They show that average grades awarded at these colleges had increased from 2.38 in 1962-3 to 2.91 (on a typical 4.0 scale) in 1985-6. They also find that the standard deviation of grades declined over this period: from 0.79 to 0.69 at Williams College, and from 0.91 to 0.77 at the other seven colleges. This suggests that as mean grades rise, the distribution becomes more concentrated.

Other work in the education research literature report similar findings. Juola (1976, 1980) found that average grade-point-average (GPA) rose by 0.432 points between 1960 and 1974. This trend was further corroborated by Levine and Cureton (1998) who find that the percentage of "A"s awarded increased from 7 percent to 26 percent, while "C"s decreased from 23 percent to 9 percent over the period 1969 to 1993. Kuh and Hu (1999) present the most data-intensive analysis to date. With a sample of 52000 students in two survey periods, the mid 1980s and the mid 1990s, they show that the average grade had improved from a "C" in the mid 80s to a "B/B-" by the mid 1990s. These papers and others are summarized in a report to the American Academy of Arts and Sciences by Rosovsky and Hartley (2002).

Before describing the model, it is useful to clarify the use of several key concepts and pieces of intuition. Firstly the term "grade inflation" is in fact a misnomer. Popular commentators have argued that grade inflation is in fact ***grade compression***, because unlike prices which

[1] Our paper, as with much of available empirical work, focuses solely on grade inflation in higher education.

[2] Their sample consisted of Amherst College, Duke University, Hamilton College, Haverford College, Pomona College, University of Michigan, UNC Chapel Hill, Williams College and University of Wisconsin.

can increase unboundedly, grades are bounded above. Thus if as much as forty to forty-five percent of a class get an "A" grade,[3] grade inflation would imply that the consumer of transcripts cannot discern between the truly outstanding from the mere average. This becomes a signal extraction problem for the consumer of grades. This reduction in information conveyed by the grade is more accurately described as grade compression. Nonetheless in keeping with the literature we will call it by its more popular name.

Note that as long as the grading scheme continues to differentiate between good types and bad types, there is no loss of information. In this case, inflation (or compression) is merely an issue of normalization. Hence to be precise, our usage of grade inflation implicitly assumes some form of dilution, i.e., grade inflation, for our purposes, is the phenomenon of passing off a bad student as a good one, by awarding him the *same grade* as the good student. This is the meaning that we wish the reader to keep in mind.

Secondly, some have made the analogy that grades function similarly to currency. Given this interpretation, widespread grade inflation would unduly penalize one's students if one did not follow suit. This has led to the observation that "bad grades drive out good", much as bad money drives out good money. These ideas have been floated variously, but have found their most cogent expression in The Economist magazine's March 7th 2002 issue (The Economist (2002)).[4]

We provide a simple model which formalizes and analyzes these ideas. We consider a labor market in which firms assign graduating students to good jobs or bad jobs at a given market wage. Not observing students' type, firms must rely on the student's transcript to ascertain the student's type. Firms do not wage-discriminate within job. This assumption is not particularly strong: equal employment laws prohibit discriminatory workplace practices, starting salaries are fairly uniform. Universities can observe students' types and care about their students' labor market outcomes. They may choose to grade-inflate, i.e., assign high grades to bad students, passing them off as good students, to boost their career opportunities. Because the firm does not wage-discriminate (within job), the average quality of all students who receive good grades forms their "collective reputation" in the labor market. Universities have an incentive to exploit this collective reputation by passing off a few more bad students as good students. Rampant grade inflation depresses this collective reputation. This manipulation is costless to individual universities, and is only limited by the extent that firms are still willing to believe that students with good grades are actually good on average.

[3] According to the Boston Globe (2001a), about half of all grades awarded at Harvard are "A/A-".

[4] Aside from the allusion to a Gresham's Law for grade inflation, we do not think it is relevant to push the analogy with money too far.

In this labor market, we show that there exists an equilibrium where every school grade-inflates. We call this a Gresham's law for grade inflation: that bad grades drive out good. This is because grade inflation takes on a contagious character: inflating schools mutually reinforce each other's practices. Furthermore, with some minor technical assumptions, we show that the equilibrium is unique. Our equilibrium also has the result that schools with higher reputations will grade-inflate more. The contagion effect may suggest how the momentum for grade inflation in the later 1960s and 1970s came about. We also show that grade inflation can be related to measures of inequality between good jobs and bad jobs. The greater the inequality between good and bad jobs, the greater the extent of grade inflation. In other words rising grade inflation seems to be related to the rising skills premium documented in the last three decades.

One of our most striking implications arises from our welfare analysis. We demonstrate that the social planner's optimum implies that there ought not be any grade inflation. Thus agents in a world with grade inflation are worse off than in a world without grade inflation. Furthermore, the presence of disciplined schools, who for some exogenous reason do not practice grade inflation, will curb the extent of grade inflation amongst the other (undisciplined) schools.

To examine the robustness of our findings, we extend our basic model to a n-job model. We find that there is still a unique equilibrium with grade inflation, as long as there is a technology jump between the best job and the second-best job. In the equilibrium, more renowned schools inflate more, and all schools and firms are worse off compared to the world without grade inflation. The only modification is that the equilibrium amount of grade inflation depending on the technology difference between the best job and the second-best job. The bigger this difference, the higher the degree of equilibrium grade inflation.

In our paper, a free-rider problem is at the heart of grade inflation. Therefore this paper is related to the large literature on public good provision, but with a twist. Here, the public good is the "collective reputation" (Tirole (1996)) of high grade students, and is exploited by all universities because the costs of actions by any one individual are spread across all schools. Our model is somewhat related to the optimal certification literature in industrial organization. Lizzeri (1999) analyzes the optimal information revelation problem faced by a certification intermediary. In his setup, he finds that a monopoly certifier will provide the minimal amount of information needed to facilitate efficient exchange, but increased competition among certifiers will induce more information to be revealed. The main difference is that in our model the objective of a certification intermediaries (university) is to maximize its agents' (students') total welfare, while in Lizzeri (1999) a certification intermediary tries to maximize its own profit by charging agents some fee.

There are two recent interesting papers also studying grade inflation. Ostrovsky and Schwarz (2003) (OS henceforth) studies the informativeness of transcripts given out by universities in a context of matching labor market. Specifically, universities chooses transcript structures and mapping rules (from ability to transcript) to maximize the average job market outcome of their students. They find that in all equilibria the same amount of information is disclosed, and full information is not disclosed in equilibria. However, OS explains grade compression more than it does grade inflation, since coarsening of information can take the form of either inflated grades or deflated grades.

Chan, Li and Suen (2003, CLS henceforth) construct a signaling model to explain grade inflation. In particular, the average quality of the students of a school can assume two values, which is the school's private information. In a semi-pooling equilibrium, with some positive probability a school gives more students a good grade when the average quality of its student is low. They further show that when the average qualities of students among schools are correlated, easy grades are strategic complements, and thus grade inflation is contagious. Though interesting, the results of CLS have some weaknesses. First, grade inflation occurs in their model only with some probability (mixed strategies), which is at odds with the fact that the average grades are constantly higher than those thirty years ago. Second, their model cannot explain why more renowned school inflate more. And finally, in deriving the result that grade inflation is contagious, they need to assume that the correlation among schools' average quality of students is common knowledge. But this correlation is usually hard to decipher due to many random shocks, for instance, competition among schools in recruiting students and business cycles. Therefore, the common knowledge assumption is relatively strong. But without that they cannot show that grade inflation is contagious.

Our paper is set out as follows: Section 2 presents the basic model. Section 3 analyzes the equilibrium with grade inflation. Section 4 considers the welfare implications and comparative statics of the model. Section 5 studies a generalized model, which yields essentially the same qualitative results. The final section concludes with some discussion and implications.

II Model Setup

There are three types of agents: students, universities[5] and firms. Students receive an education from universities. Upon graduation they enter the job market, and are recruited by firms to become workers. Each matched job yields output, which is split between the worker and the firm. We describe each agent in more detail below.

[5] we sometimes use "school" and "university" interchangeably.

Students (workers) Students may differ along two dimensions: type, and transcript. We assume students are of two ability types $t \in \{H, L\}$ (H denotes high ability and L denotes low ability), while their transcript (or grades) might read as $\tau \in \{h, l\}$ (h denotes high grades and l denotes low grades), which may be different from their type t. Note that upper case denotes actual type, but lower case denotes transcript. It is easiest to think of students as graduating college seniors applying for jobs. The students are enrolled at the universities described below.

Universities (or Schools) Universities are indexed by $i \in \{1, 2, ..., n\}$ where n is large. Each university enrolls a proportion π_i of all students, the total population of which satisfies $\sum_i \pi_i = 1$. We assume that every university has a small enrolment relative to the entire cohort of students (i.e., π_i is small relative to 1). Let $\lambda_i \in (0, 1)$ denote the exogenously given proportion of H types enrolled in university i. The parameter λ_i is our measure of the reputation of university i, and without loss of generality, let $\lambda_n > \lambda_{n-1} > ... > \lambda_1$ so that larger n corresponds to a more reputable/better university. We assume that (1) a student's type can be observed by his university (such as from student-teacher interaction, homework and class participation), and (2) universities always award grade h to students whose type is H (bright students always do well). Universities may also be tempted to award grade h to students whose type is L, possibly to improve their career opportunities. Specifically, university i may choose to award a random proportion θ_i of type L students, a grade of h. A choice of $\theta_i = 1$ implies maximal inflation: all L students are awarded h transcripts and passed off as H types. If $\theta_i = 0$, the university reports truthfully each student's type.

Each university cares about the total expected wage payment of its graduates and the proportion of its graduates assigned to good jobs. The university receives $\varepsilon > 0$ additional utility for each of its students assigned to the good job.[6] ε captures the idea of a prestige premium for good jobs (good placements reflect well on the university). More specifically, each university's objective is to maximize the total expected wage payment of its graduates plus its total prestige premium. We consider this to be a reasonable assumption because universities care about the labor market success of their students at least indirectly, and consequently often report mean starting salaries and average employment outcomes of their graduates in their promotional material. These indicators feature prominently in university quality and reputation rankings, which they ultimately care about.[7]

[6] This assumption is kind of technical. We assume this to pin down the unique equilibrium. The results of our paper still hold when ε is very small.

[7] This alludes to the so-called trend of commercialization in higher education, observed since the later 60s and early 70s.

Firms We assume that there are many firms, each of which has two types of vacancies. There are two kinds of jobs: $j \in \{g, b\}$, where g denotes good job and b denotes bad job. Let output be denoted as y, which depends on job type j and the type of worker t. In particular,

$$y(H, g) = 1$$
$$y(L, g) = 0$$
$$y(H, b) = y(L, b) = \alpha, \text{ where } 0 < \alpha < 1$$

These assumptions capture the feature that good jobs need to be matched with high type workers in order to produce an output of 1. Otherwise, output would be 0. Worker misassignment in good jobs is costly to the firm. Unlike good jobs, bad jobs are not sensitive to worker quality and produce a moderate output $\alpha \in (0, 1)$ regardless of worker type. Clearly, if there is great uncertainty over a worker's type, the firm may be better off assigning that worker to the bad job, as it yields α for sure.

Firms cannot observe students' (workers') type, but can observe students' transcripts (grades), and their alma mater's reputation λ_i. We further assume that firms also observe θ_i, each school's degree of grade inflation.[8] The firm's problem is to assign the applicant to either the good job or the bad job, at the prevailing wage, which is determined in the market. We assume that the firms face a non-discrimination restriction: all students assigned to the same job have to be paid the same wage, even if they were from different colleges. This is an important assumption. We see this as reasonable upon the following grounds: First, as pointed out by Stephen Coate and Glenn C. Loury (1993), "[d]iscriminatory wages for the same work is a flagrant violation of equal-employment laws". Secondly, Andrew Weiss (1980) presents some empirical evidence that firms pay workers a more or less uniform wage rate despite large productivity differences between workers. Weiss argued that this was due to union pressure and the adverse morale effects and/or hostile work environments which arise as a result of unequal wages.

Given this assumption, firms must make their wage payment w, upon job characteristics: that is, firms either pay w_g to workers assigned to good jobs or pay w_b to workers matched with bad jobs.[9] We also assume that firms are risk neutral.

[8] This is equivalent to assuming that firms know the proportion of students of school i receiving grades h. Combining this with knowledge of λ_i, they can infer θ_i. We make this assumption because we want to show that even if firms have information on the grade distribution of each school, grade inflation still remains. This has important policy implications in light of practices by some universities reporting grade distributions on students' transcripts. In the context of our model we can show this has no effect.

[9] As a robustness check, we show in a later section that this assumption can be modified and grade inflation will still result.

Firms assign jobs by maximizing the expected profit from each job, taking wages w_j as given. Formally,

$$\max_j Ey\left(t,j\right) - w_j$$

Wage Determination We assume wages are determined according to a Nash bargaining rule in the market. Let $\gamma \in (0,1)$ be the firm's bargaining power, $1-\gamma$ be the worker's bargaining power,[10] and let the outside options of both firms and workers be normalized to zero. Let A be the set of schools whose h-grade students will be assigned to job g by firms (acceptance set). A is endogenous and will be determined in equilibrium. For good jobs, the average expected total surplus (ATS_g) is

$$ATS_g = \frac{\sum_{i\in A}\pi_i E\left[y\left(t|h_i,g\right)\right]}{\sum_{i\in A}\pi_i}$$

Therefore, $w_g = (1-\gamma)ATS_g$. For bad jobs, the average expected total surplus $ATS_b = \alpha$, independent of worker types, thus $w_b = (1-\gamma)\alpha$.

We now formally write the universities' objective function. Taking λ_i, w_j as given, university i chooses θ_i such that

$$\max_{\theta_i}[\lambda_i + (1-\lambda_i)\,\theta_i](w_g+\varepsilon) + (1-\lambda_i)(1-\theta_i)w_b \tag{1}$$

$[\lambda_i + (1-\lambda_i)\,\theta_i]$ is the total proportion of h-grade students in university i. These students have some hope of being assigned to good jobs and getting w_g. The remaining proportion of students $(1-\lambda_i)(1-\theta_i)$ receive grade l. Since H workers do not get l, the true type L of l-transcript workers is revealed to firms. Thus they will be assigned to bad jobs for sure, and get w_b.[11]

III The Equilibrium with Grade Inflation

Although there are three types of agents in the model, only two types are active agents: universities and firms; students are passive in the sense that do not take actions that affect their own payoff. Each university's strategy is to choose $\theta_i \in [0,1]$; and each firm's strategy is to assign a job to each student.

[10] The bargining power parameter γ does not affect the equilibrium of the model. γ can be very close to zero (thus the job market is very close to perfectly competitive) and the equilibrium of the model is still the same. We make this technical assumption to get rid of multiplicity of equilibria by making firms care about the job assignment.

[11] It is without loss of generality to put the same weight on total expected wage payment and prestige premium in a university's objective function. This is simply a normalization of ε.

The timing of the game is the following: first all schools choose θ_i simultaneously; then observing θ_i, firms make job assignment decisions taking market wages as given; thereafter production takes place. The acceptance set A (firms' job assignment decision) and market wages are determined simultaneously. This is a dynamic game with incomplete information. We adopt perfect Bayesian Nash (PBE) equilibrium as our solution concept.

Definition 1 *The Perfect Bayesian Nash equilibrium of the game is the following:*

1. *given other schools' strategies and firms' job assignment strategy, each school i chooses θ_i to maximize the total wage payment plus prestige premium of its students;*

2. *given firms' belief about students' quality, they assign jobs to maximize their own profit;*

3. *firms' beliefs are derived according to Bayes' Rule.*

A Firms

The firm knows all l-transcript students from all schools are of type L, and will assign them to job b. The total surplus of this match is α, so the wage payment is $w_b = (1-\gamma)\alpha$ and the firm's profit is $\gamma\alpha$. Payoffs are identical when type H workers are matched to bad jobs. On the other hand, when a h-transcript applicant from university i is hired by a firm, the firm, not fully believing the transcript it sees, needs to assess the student's probability of being type H. Let $\Pr(H|h_i)$ be the firm's posterior belief that the h-grade student from university i is type H. By Bayes' rule,

$$\Pr(H|h_i) = \frac{\lambda_i}{\lambda_i + (1-\lambda_i)\,\theta_i} \equiv f_i(\theta_i)$$

Clearly, $f_i'(\theta_i) < 0$, and $f_i(\theta_i) \in [\lambda_i, 1]$ with $f_i(0) = 1$ and $f_i(1) = \lambda_i$. The greater the extent of grade inflation by university i, the less the firm will believe that this worker is a type H.

The algorithm for determining the firms' acceptance set A is as follows. Given each school's strategy θ_i and reputation λ_i, firms will form beliefs $f_i(\theta_i)$ about each university and can rank each school accordingly: the higher the $f_i(\theta_i)$, the higher the expected quality of the h-grade students. Denote the school with highest $f(\theta)$ as r_1, the school with second highest $f(\theta)$ as r_2 etc. Firms make job assignment decisions according to $f_i(\theta_i)$. Given firms' job assignment decision, a candidate acceptance set A will be determined. Given this candidate A, the market wages w_g and w_b are determined. For A to be an equilibrium acceptance set, it should be consistent with firms' job assignment strategy: i.e., given A, firms' job assignment decision should be optimal.

We start with A being the set of all schools and proceed by elimination. First suppose firms assess all schools to belong to A. Under this candidate A,

$$w_g = (1-\gamma)\sum_i \pi_i f_i(\theta_i)$$

The consistency requirement is satisfied if and only if firms are willing to assign h-grade students from school r_n (the lowest ranked school) to job g. That is,

$$f_{r_n}(\theta_{r_n}) - (1-\gamma)\sum_i \pi_i f_i(\theta_i) \geq \gamma\alpha \tag{2}$$

The left hand side of (2) is the firm's profit if it assigns the student to a job g: $f_{r_n}(\theta_{r_n})$ is the expected output, the second term is the wage. The right hand side is the firm's profit if he assigns the student to job b. If (2) is satisfied, then the equilibrium acceptance set A includes all schools. Otherwise, school r_n is excluded from the acceptance set A.[12]

Next firms consider the set of all schools except school r_n as the candidate A. Consistency requires

$$f_{r_{n-1}}(\theta_{r_{n-1}}) - (1-\gamma)\frac{\sum_{r_i<r_n} \pi_i f_i(\theta_i)}{\sum_{r_i<r_n} \pi_i} \geq \gamma\alpha \tag{3}$$

If (3) is satisfied, then the equilibrium A includes all the schools except school r_n. Otherwise, school r_{n-1} is excluded from A. Firms next consider the set of all schools except r_n and r_{n-1} as the candidate acceptance set. This process will go on until they reach a consistent acceptance set A, which will be the equilibrium set A.

Another way to think about the determination of the equilibrium set A is the following. There is a hypothetical Walrasian auctioneer in the market who announces a candidate acceptance set A and the corresponding market wage w_g. Firms then announce the acceptance set A given w_g. If the two announcements of A coincide, consistency is satisfied. There may be multiple consistent acceptance sets, but our algorithm picks the largest consistent set as the equilibrium set A. This is reasonable since competition among firms would drive the equilibrium set A to be the largest consistent A.

Given firms' optimal job-assignment strategy, it is easy to check that $w_g \geq w_b$; otherwise firms would assign all h students to job b.[13] Now the necessary condition for school i's h-grade students to be assigned to job g is

$$f_i(\theta_i) \geq \alpha \tag{4}$$

[12] If (2) is not satisfied, school r_n is excluded from the equilibrium set A. Since r_n is the lowest-ranked school, w_g increases as A shrinks.

[13] If $w_g < w_b$ it must be the case that all school over-inflated.

The proof is straightforward. Suppose not, $f_i(\theta_i) < \alpha$. Note that $w_g \geq w_b = (1-\gamma)\alpha$. Therefore,

$$f_i(\theta_i) - w_g \leq f_i(\theta_i) - (1-\gamma)\alpha < \gamma\alpha \tag{5}$$

and all the h-grade students from school i will be assigned to job b.

B Universities

Having derived firms' optimal job assignment strategy given their beliefs, we turn to universities' optimal strategy. Formally, given other universities' strategies and firms' strategy, each school i chooses θ_i

$$\max_{\theta_i}[\lambda_i + (1-\lambda_i)\,\theta_i][(1-\gamma)\frac{\sum_{x\in A}\pi_x f_x(\theta_x)}{\sum_{x\in A}\pi_x} + \varepsilon] + (1-\lambda_i)(1-\theta_i)(1-\gamma)\alpha \tag{6}$$

$$\text{subject to} \;\; f_i(\theta_i) - (1-\gamma)\frac{\sum_{x\in A}\pi_x f_x(\theta_x)}{\sum_{x\in A}\pi_x} \geq \gamma\alpha \tag{7}$$

Condition (7) makes sure that firms are willing to assign h-grade students from school i to job g. To derive the Nash equilibria, we prove a series of lemmas.

Lemma 1 *If all the other schools $k \neq i$ do not grade-inflate ($\theta_k = 0$), then school i always has an incentive to grade-inflate, and it is optimal for it to do so until $f_i(\theta_i) - (1-\gamma) = \gamma\alpha$.*

Proof. *Given that no other schools inflate, the programming problem of school i is:*

$$\max_{\theta_i}[\lambda_i + (1-\lambda_i)\,\theta_i][(1-\gamma) + \varepsilon] + (1-\lambda_i)(1-\theta_i)(1-\gamma)\alpha \tag{8}$$

$$\textit{subject to } f_i(\theta_i) - (1-\gamma) \geq \gamma\alpha \tag{9}$$

Here we use the fact that each school's student population π_i is small, so one school's inflation has a negligible effect on w_g. Ignoring constraint (9), (8) *is obviously increasing in θ_i since $\alpha < 1$. So the school has an incentive to inflate as much as possible. But this school is disciplined by the firms' constraint* (9)*: that too much inflation will result in their h-grade students being assigned to job b. So the school will choose the maximum θ_i^* such that* (9) *binds, i.e., $f_i(\theta_i^*) - (1-\gamma) = \gamma\alpha$. Since $f_i(\theta_i^*) = (1-\gamma) + \gamma\alpha < 1$, $\theta_i^* > 0$.* ■

The intuition for this result is the following. If other schools do not inflate, the market wage for job g is relatively high. Since school i's population is small, its grade inflation will not affect the market wage. Because firms cannot wage-discriminate among workers assigned to the same job but are from different schools, school i has an incentive to maximize the number of its students being assigned to job g. Thus grade inflation results. However, the school cannot grade-inflate too much because then firms would assign all its students to job b. So the school will inflate until firms are indifferent between assigning its h-grade students to job g and job b, given by the condition $f_i(\theta_i^*) = (1-\gamma) + \gamma\alpha$.

From lemma 1, it should emerge that a sort of free-rider problem is at the heart of grade inflation. The market wage w_g is determined by the average posterior of h-grade students of all the schools in acceptance set A. The average quality of h-grade students in the acceptance set A in some sense forms their "collective reputation" in the labor market, which is a kind of public good. Given that each school's population is small, each school has incentive to free ride (grade-inflate), without considering its negative impact on the public good.

Lemma 2 *The more other schools inflate their grades, the more the school i will inflate its grades.*

Proof. *Suppose that all other schools $k \neq i$ inflate to the extent such that $f_k(\theta_k) = \overline{f} \in (\alpha, 1)$. Then* (6) *and* (7) *can be simplified as*

$$\max_{\theta_i}[\lambda_i + (1-\lambda_i)\,\theta_i][(1-\gamma)\overline{f} + \varepsilon] + (1-\lambda_i)(1-\theta_i)(1-\gamma)\alpha \tag{10}$$

$$\textit{subject to } f_i(\theta_i) - (1-\gamma)\overline{f} \geq \gamma\alpha \tag{11}$$

Since $\overline{f} > \alpha$, school i would choose to maximize θ_i until constraint (11) *binds. The optimal solution θ_i^* therefore satisfies $f_i(\theta_i^*) = (1-\gamma)\overline{f} + \gamma\alpha$. Note that $f_i(\theta_i^*)$ is increasing in $\overline{f}$, which means that θ_i^* is decreasing in $\overline{f}$. This implies that the more other schools inflate (lower $\overline{f}$), the more school i will inflate.* ■

This result shows that grade inflation is mutually encouraging or contagious. The intuition is the following. The more other schools inflate (of course not too much), the lower the market wage for job g. Given that $w_g \geq w_b$, school i always wants to inflate as much as possible subject to the constraint that firms are still willing to assign its h students to job g. Firms will indeed be more willing to do so, since the market wage is decreasing with the extent of other schools' grade inflation. Hence firms' increased tolerance of inflation encourages school i to inflate more.

Although the result of lemma 2 is not from equilibrium analysis (it is a kind of best response function), we believe that it has a relevant interpretation. Consider a dynamic interpretation of the Cournot model. Suppose that the schools are divided into two sets 1 and 2 with equal population 1/2, and the game is played repeatedly. In each odd period, schools in set 1 choose θ_i, and in each even period schools in set 2 chooses θ_k. In the starting period all schools do not grade-inflate. In this dynamic setting, θ_i and θ_k will chase each other over time, until they reach the equilibrium. Grade inflation is clearly mutually encouraging and contagious in this setting. We summarize this result in proposition 1.

Proposition 1 *Grade inflation is contagious.*

C The Equilibrium

Lemma 3 *It is never optimal for a school to inflate too much such that all its students are assigned to job b.*

Proof. *If a school, say i, inflates too much such that all its students are assigned to job b, its payoff per student is $(1-\gamma)\alpha$. If, instead, school i does not inflate ($\theta_i = 0$ and $f_i(\theta_i) = 1$), its h-grade students will be assigned to job g for sure, because $1-(1-\gamma) > \gamma\alpha$, where $(1-\gamma)$ is the highest possible wage. In this case, school i's average payoff per student is $\lambda_i(w_g+\varepsilon)+(1-\lambda_i)(1-\gamma)\alpha$. Given $w_g \geq w_b = (1-\gamma)\alpha$,*

$$\lambda_i(w_g+\varepsilon)+(1-\lambda_i)(1-\gamma)\alpha > (1-\gamma)\alpha$$

■

In other words, inflating too much such that all its students are assigned to job b is dominated by the strategy of no inflation at all. Lemma 3 tells us that to find an equilibrium, we can restrict attention to the strategy profiles such that h_i, $\forall\ i$, are assigned to job g; i.e., the h-grade students of all universities are assigned to good jobs. Put in another way, without loss of generality we can restrict our attention to the case where the acceptance set A includes all schools. As we argued before, the necessary condition for school i's h-grade students to be assigned to job g is that $f_i(\theta_i) \geq \alpha$. So we can restrict attention to the strategy profiles such that $f_i(\theta_i) \geq \alpha$ for all i.

Lemma 4 *It is never optimal for school i to choose $\theta_i < 1$ such that $f_i(\theta_i) > f_k(\theta_k)$ for some school $k \neq i$.*

Proof. *From Lemma 3 we know that we can restrict attention the strategy profiles such that h-grade students from all schools are assigned to job g. Consider a strategy profile such that school i inflates less than at least one other school, say k. i.e., $f_i(\theta_i) > f_k(\theta_k)$. This cannot be optimal because school i can still inflate more to θ_i' such that $f_i(\theta_i') = f_k(\theta_k)$. Hence by increasing inflation to θ_i', more students from i will be assigned to job g, and its payoff increases while w_g is unchanged (h-grade students from k are assigned to g by our initial assumption) Therefore, a strategy profile with $f_i(\theta_i) > f_k(\theta_k)$ cannot be a Nash equilibrium. The only exception is if school i has already inflated to its maximum amount (giving all its students grades h), so they can not inflate more.* ■

Lemma 4 implies that there are two types of candidate equilibria: (i) all schools inflate until $f_i(\theta_i) = f_k(\theta_k)$ for all i and k (interior solution) and all $\theta_i < 1$; (2) all schools i inflate until $f_i(\theta_i) = f_k(\theta_k)$ (for all i and k) for $\theta_i < 1$, and other schools x will inflate until $\theta_x = 1$ (corner solution).

Proposition 2 *(Gresham's Law of Grade Inflation) There is a unique equilibrium. If $\lambda_n \leq \alpha$, then in equilibrium all schools inflate until $f_i(\theta_i^*) = \alpha$. If $\lambda_{n-1} \leq \alpha < \lambda_n$, then in equilibrium, $\theta_n^* = 1$, and other schools inflate to the extent such that $f_i(\theta_i) = f_k(\theta_k) = \alpha$.*

Proof. *By Lemma 4, we can restrict attention to two types of candidate equilibrium. We first look at the case involving an interior solution. Suppose that all schools inflate such that $f_k(\theta_k) > \alpha$, $\forall\ k$. Now consider school i's incentive given other schools' $f_k(\theta_k)$.* (6) *and* (7) *can be simplified as*

$$\max_{\theta_i}[\lambda_i + (1-\lambda_i)\,\theta_i][(1-\gamma)f_k(\theta_k) + \varepsilon] + (1-\lambda_i)(1-\theta_i)(1-\gamma)\alpha \tag{12}$$

$$\textit{subject to } f_i(\theta_i) - (1-\gamma)f_k(\theta_k) \geq \gamma\alpha \tag{13}$$

By constraint (13)*, school i will inflate to θ_i' such that $f_i(\theta_i') = (1-\gamma)f_k(\theta_k) + \gamma\alpha < f_k(\theta_k)$, and firms are still willing to assign its h-grade students to job g. This implies that $\theta_i' > \theta_i$. Since $f_k(\theta_k) > \alpha$, by* (12) *school i's total payoff is strictly increasing in θ_i. Thus, school i has incentive to increase its degree of inflation to θ_i'. Therefore, This strategy profile cannot be a Nash equilibrium. Now the only candidate left is the strategy profile where all schools inflate until $f_k(\theta_k) = \alpha$, $\forall\ k$. Given other schools' strategies $f_k(\theta_k)$, by* (12) *school i strictly prefers grade-inflating to θ_i^* such that $f_i(\theta_i^*) = \alpha$; although w_g is independent of θ_i, school i gets maximum prestige premium by inflating to θ_i^* . Therefore, this strategy profile does constitute a Nash equilibrium. Moreover, it is unique. Note that this equilibrium exists if and only if λ_n, the most reputable school's reputation, is less than or equal to α. Otherwise, at least $f_n(\theta_n) > \alpha$, which is inconsistent with the equilibrium condition.*

The case $\lambda_{n-1} \leq \alpha < \lambda_n$ is very similar to the case $\lambda_n \leq \alpha$. The strategy "all the schools inflate until $f_k(\theta_k) > \lambda_n$, $\forall\ k$" is clearly not an equilibrium, since all schools have incentive and freedom to inflate more given other schools' strategies. The strategy profile such that $\theta_n = 1$ and all other schools inflate until $\lambda_n > f_k(\theta_k) > \alpha$ is not an equilibrium either Because school n (which is maximally inflated) cannot affect the market wage for job g individually, school i's programming problem is still captured by (12) *and* (13)*, and all other firms can still inflate more. So the only Nash equilibrium is that all other firms inflate until $f_k(\theta_k^*) = \alpha$ for $k \neq n$ and $\theta_n^* = 1$.* ■

For the case where the reputation of several schools exceeds α, the exact characterization of the equilibrium is bit more complicated but the basic features of the equilibrium are the same: the most reputable schools inflate to their maximum, and all the other schools inflate until $f_k(\theta_k) = f_i(\theta_i) > \alpha$ for all i and k. For the sake of easy exposition, in the rest of the paper we assume that $\lambda_n \leq \alpha$. In other words, we restrict our attention to the interior solution case.

Proposition 2 is a very strong result. It implies that in the unique equilibrium, *all* schools will inflate until $f_i(\theta_i) = \alpha$ for all i. In equilibrium, firms are indifferent between assigning h-grade students to jobs b and g ((13) is binding). This is because each school ignores the negative externality it imposes on the collective reputation of h-grade students in the labor market. Given that the market wage for good jobs only depends on the collective reputation, each school has an incentive to exploit this potential free ride by grade inflation, which gives its students a career boost. However if all the schools grade-inflate, the collective reputation in the labor market is substantially "milked". As a result, the market wage for good jobs gets pushed downwards. This is very similar to the classical result of the under-provision of a public good. In our model, collective reputation is the public good, and each school's self discipline (not to grade-inflate) somewhat represents the individual contribution to the public good.

Following The Economist magazine, we have depicted a "Gresham's law" for grade inflation: that "bad" grades drive out "good". Indeed, the synonym is remarkably apt. All schools will inflate in equilibrium, thus rendering their grades unreliable, (hence bad), while truthful (good) grades are a rare sight because individual universities see no gain in bucking the grade inflation trend.

Note that adverse selection plays an important role in our model. Should firms observe the true type of each student, they could easily assign him according to his true type. Market wages would then be independent of the strategies of schools. It is the combination of adverse selection and firms' non-discrimination that causes the free-rider problem: the market wage for good jobs depends on the collective reputation of all h-grade students in the labor market.

IV Welfare Implications and Comparative Statics

A Welfare analysis

Define the total social surplus SS as the expected output from all the matched jobs.[14] Formally,

$$SS = \sum_i \pi_i[\lambda_i \times 1 + (1-\lambda_i)\theta_i \times 0 + (1-\lambda_i)(1-\theta_i) \times \alpha] \tag{14}$$

From (14) it is obvious that SS is maximized when $\theta_i = 0$ for all i. A social planner can achieve this first-best outcome by forbidding the practice of grade inflation (if at all possible). But the equilibrium outcome of the models result in positive grade inflation: $f_i(\theta_i^*) = \alpha < 1 \Rightarrow \theta_i^* > 0$. This reduces the social surplus. This welfare loss comes from job mis-assignment: due to grade inflation, some L workers are mis-assigned to job g resulting in an output of 0, when instead, they should be assigned to job b, resulting in an output of α.

Firms are also worse off in equilibrium compared to the first-best outcome (no grade inflation). In equilibrium, their expected profit from a job g is $\gamma\alpha$, while in the first-best outcome their expected profit from a job g is γ.

A stronger result is that schools are also worse off in equilibrium relative to the first-best outcome. Without grade inflation, school i's total payoff is $TW_i(NI)$[15] (NI denotes no-inflation, and the school i sums payoff over π_i students)

$$\begin{aligned} TW_i(NI) &= \pi_i\{\lambda_i \times [(1-\gamma)+\varepsilon] + (1-\lambda_i) \times (1-\gamma)\alpha\} \\ &\approx (1-\gamma)\pi_i[\lambda_i + (1-\lambda_i)\alpha] \end{aligned}$$

In the equilibrium outcome, school i's total payoff is $TW_i(EQ)$ (EQ denote equilibrium)

$$TW_i(EQ) = \pi_i\{[\lambda_i + (1-\lambda_i)\theta_i] \times [(1-\gamma)\alpha + \varepsilon] + (1-\lambda_i)(1-\theta_i) \times (1-\gamma)\alpha\} \tag{15}$$

$$\approx (1-\gamma)\pi_i\alpha \tag{16}$$

Clearly, $TW_i(NI) > TW_i(EQ)$. That is, every school is worse off in the grade inflation equilibrium. In fact, all schools collectively share a $(1-\gamma)$ portion of output loss from job mis-assignment. But the problem is that each individual school does not bear the cost of

[14] We exclude the prestige premium of good jobs enjoyed by schools from the social surplus. Recall that by assumption, they are small relative to α.

[15] Here we use the assumption that the prestige premium for good job, ε is small.

its grade inflation since the wage payment only depends on the collective reputation of all schools. If all schools inflate, the collective reputation depreciates substantially and each school is worse off. The central problem is that no individual school has incentive to maintain the collective reputation, but instead want to free-ride at the expense of other schools and firms. The following proposition summarizes the welfare analysis.

Proposition 3 *Compared to the first-best outcome (no grade inflation), in equilibrium all schools and firms are worse off, and society as a whole is also worse off.*

B Comparative statics

We now turn to comparative statics. In equilibrium, all schools inflate their grades such that $f_i(\theta_i^*) = \alpha$. This implies that

$$\theta_i^* = \frac{\lambda_i(1-\alpha)}{(1-\lambda_i)\alpha} \; \forall i \tag{17}$$

From (17), we observe that θ_i^* is increasing in λ_i. This means that schools with a greater reputation will inflate more than schools with less reputation. The intuition for this result is that a larger endowment of good students gives the school greater leeway to grade-inflate since it would not decrease firms' posterior beliefs $f(\theta)$ by much.

From (17), it is also obvious that θ_i^* is decreasing in α. The parameter α measures the productivity differential between good jobs and bad jobs. The larger the job productivity (equivalently skills) differential, the higher the grade inflation. The intuition for this result is the following: the larger the job productivity differential, the more a good job will be attractive relative to a bad job. This increases schools' incentive to grade-inflate. On the other hand, firms are more willing to tolerate grade inflation by assigning h-grade students to good jobs since a good job is far more productive if it matches with a H worker. Put another way, because bad jobs are really bad, the school would rather pass off bad students as good in the hope that some of them could secure good jobs. Thus, we have the following proposition

Proposition 4 *Other things equal, in equilibrium more renowned schools will inflate more than less renowned schools; and the bigger the productivity differential between good job and bad job, the higher the degree of grade inflation for all the schools in equilibrium.*

In the baseline model, we assume that all schools behave in an opportunistic way. Now suppose that there is a subset of schools D (disciplined) with a total population of students $\pi_D >> 0$, who for some exogenous reasons eschew grade inflation.[16] The rest of schools,

[16] These schools may have a different objective function.

denoted as set ND (non-disciplined) still behave in an opportunistic way, with a total population of students $1-\pi_D$. In this variation of the model, the equilibrium is characterized by the following lemma.

Lemma 5 *In the equilibrium, for any school $k \in ND$, it will grade-inflate until*

$$f_k(\theta_k^*) - (1-\gamma)[\pi_D + (1-\pi_D)\sum_{k\in ND}\pi_k f_k(\theta_k^*)] = \gamma\alpha \tag{18}$$

is satisfied.

Proof. *By a similar argument to the baseline model, in any candidate equilibrium, any school $k, m \in ND$ will inflate until $f_k(\theta_k) = f_m(\theta_m)$. Moreover, the equilibrium is unique and for any school $k \in ND$, firms are indifferent between assigning its h-grade students to job g and job b. That is:*

$$f_k(\theta_k) - w_g = \gamma\alpha$$

where the equilibrium market wage w_g is:

$$w_g = (1-\gamma)[\pi_D + (1-\pi_D)\sum_{k\in ND}\pi_k f_k(\theta_k)]$$

Solving the above two equations, the equilibrium grade inflation θ_k^* for $k \in ND$ is characterized by

$$f_k(\theta_k^*) - (1-\gamma)[\pi_D + (1-\pi_D)\sum_{k\in ND}\pi_k f_k(\theta_k^*)] = \gamma\alpha \tag{19}$$

■

It is obvious that $f_k(\theta_k^*) > \alpha$. That is, the equilibrium grade inflation θ_k^* is less than when $\pi_D = 0$. The existence of disciplined schools makes non-disciplined schools inflate less in the equilibrium. This is because disciplined schools push up the market wage w_g. Although non-disciplined schools now have more incentive to grade-inflate, they cannot be too outrageous because given a higher w_g, the cutoff value of the posterior belief whereby firms are indifferent between assigning h-grade students from non-disciplined schools to good jobs and bad jobs is also higher. This forces non-disciplined schools to inflate less in equilibrium. Thus disciplined schools discipline the non-disciplined schools.

From equation (19), it is easy to check that $f_k(\theta_k^*)$ is increasing in π_D. This means that as the set of disciplined schools expands, non-disciplined schools inflate less in the equilibrium.

Compared to the case without disciplined schools, all agents are better of in the world with some discipline. Less job mis-allocation means less loss in total surplus, making firms better off. Both disciplined and non-disciplined schools are better off due to a higher equilibrium market wage w_g. Of course, the non-disciplined schools gains more than the disciplined schools. The following proposition summarizes the results of this comparative statics analysis.

Proposition 5 *The existence of disciplined schools forces non-disciplined schools to grade-inflate less in equilibrium. The larger the proportion of disciplined schools, the smaller the extent of grade inflation by non-disciplined schools. Compared to the case without disciplined schools, all the agents are better off (in equilibrium) with the presence of disciplined schools.*

V Extensions

In the basic model, we assume that there are only two kinds of jobs available. In this section, we introduce another type of job m (intermediate), besides job g and job b. The technology of job m is the following:

$$y(H,m) = \overline{\alpha};\ y(B,m) = \underline{\alpha};\ 0 < \underline{\alpha} < \alpha < \overline{\alpha} < 1.$$

That is, job m is less productive than job g but more productive than job b if all are matched with a H type worker. On the other hand, job m is more risky than job b but safer than job g if all are matched with a L type worker. We further assume that

$$\frac{\underline{\alpha}}{1-\overline{\alpha}+\underline{\alpha}} > \alpha \tag{20}$$

Assumption (20) makes sure that job m is not dominated: there is a range of f (the probability that a student with grade h is a of H type) such that assigning a h student to job m is optimal. Given (20), the socially optimal job assignment is the following: there are $\widehat{f_g} = \frac{\underline{\alpha}}{1-\overline{\alpha}+\underline{\alpha}}$ and $\widehat{f_m} = \frac{\alpha-\underline{\alpha}}{\overline{\alpha}-\underline{\alpha}}$ ($\widehat{f_g} > \widehat{f_m}$) such that if $f \in [\widehat{f_g}, 1]$ assigning a h student to job g is optimal, if $f \in [\widehat{f_m}, \widehat{f_g}]$ assigning a h student to job m is optimal and if $f \in [0, \widehat{f_m}]$ assigning a h student to job b is optimal. The technology of all the job assignments are illustrated in figure 1.

To simplify analysis, we assume that $\lambda_n < \widehat{f_g}$ to avoid corner solutions. First, we identify a candidate equilibrium. The strategy profile is the following: each school i inflate to the extent such that $f_i(\theta_i) = \widehat{f_g}$ and all their h student are assigned to job g. In this candidate equilibrium, $w_g = (1-\gamma)\widehat{f_g}$, $w_b = (1-\gamma)\alpha$ and w_m is not determined. Next, we prove that this indeed is an equilibrium. There are two kinds of possible deviations for school i: deviate to some $\overline{f} > \widehat{f_g}$ or deviate to some $\underline{f} \in [\widehat{f_m}, \widehat{f_g})$ (note that deviating to some $f < \widehat{f_m}$ is obviously not optimal). Consider a deviation to $\overline{f} > \widehat{f_g}$ (inflate less). The h students of school i are still

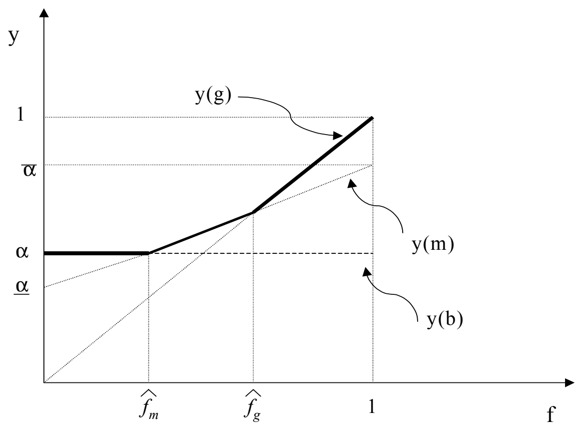

Figure 1:

assigned to job g, and since the population of each school is relatively very small, w_g is not affected. But now less student are assigned to job g, the school's payoff is strictly decreased since $w_g > w_b$. Next consider a deviation to $\underline{f} \in [\widehat{f_m}, \widehat{f}_g)$ (inflate more). Now the h students of school i are assigned to job m and get a wage $w_m = \overline{\alpha}\underline{f} + \underline{\alpha}(1-\underline{f})$ (note that we are considering school i's unilateral deviation). School i's payoff following the equilibrium strategy is:

$$\frac{\lambda_i}{\widehat{f}_g}(1-\gamma)\widehat{f}_g + (1-\frac{\lambda_i}{\widehat{f}_g})(1-\gamma)\alpha = (1-\gamma)[\alpha + \lambda_i(1-\frac{\alpha}{\underline{\alpha}}(1-\overline{\alpha}+\underline{\alpha}))] \tag{21}$$

School i's payoff if he deviates to $\underline{f}$ is

$$\frac{\lambda_i}{\underline{f}}(1-\gamma)[\overline{\alpha}\underline{f} + \underline{\alpha}(1-\underline{f})] + (1-\frac{\lambda_i}{\underline{f}})(1-\gamma)\alpha \tag{22}$$

(21)-(22) equals to

$$\begin{aligned}
& (1-\gamma)\lambda_i\{1-\frac{\alpha}{\underline{\alpha}}(1-\overline{\alpha}+\underline{\alpha}) - \overline{\alpha} + \underline{\alpha} - \frac{1}{\underline{f}}(\underline{\alpha}-\alpha)\} \\
= \; & (1-\gamma)\lambda_i(\frac{\alpha}{\underline{\alpha}}-1)(\overline{\alpha}-\underline{\alpha}-1+\frac{\alpha}{\underline{f}}) \\
> \; & (1-\gamma)\lambda_i(\frac{\alpha}{\underline{\alpha}}-1)(\overline{\alpha}-\underline{\alpha}-1+\frac{\alpha}{\widehat{f}_g}) = 0
\end{aligned}$$

Therefore, school i has no incentive to deviate to any $\underline{f} \in [\widehat{f_m}, \widehat{f}_g)$. This property can be understood in the following intuitive way. Compared to inflating to $\widehat{f}_g$, inflating to any $\underline{f} \in$

$[\widehat{f_m}, \widehat{f_g})$ is socially inefficient. Since all the other schools inflate to $\widehat{f_g}$, if deviates to $\underline{f}$, school i's total payoff from its student is proportional to the total social surplus that they created. As a result, inflating to $\underline{f}$ is dominated by inflating to $\widehat{f_g}$.

Thus, all schools inflating to $\widehat{f_g}$ is an equilibrium. Moreover, it is unique. We prove this in several steps. First, in any equilibria all schools should inflate to the same f (ignoring the case of corner solution). The reasoning is very similar to those of lemma 4: if this is not the case, one school who inflate less can inflate more, still have their all h student assigned to the same job and getting the same wage as before. Second, suppose all schools inflate to some $\overline{f} > \widehat{f_g}$. Then one school would have incentive to inflate a little bit more: its h student are still assigned to job g and getting the same wage as before. Third, suppose all schools inflate to some $\underline{f} \in [\widehat{f_m}, \widehat{f_g})$. Then all h students are assigned to job m. But now one school would have incentive to inflate to $\widehat{f_g}$ and all its h student are assigned to job g and get a wage $(1-\gamma)\widehat{f_g}$. This deviation yields a higher payoff (by the previous proof (21)-(22)> 0). Finally, all school inflating to any $f < \widehat{f_m}$ is clearly not an equilibrium.

To sum up, the game has a unique equilibrium: all schools inflate to the extent such that $f_i(\theta_i) = \widehat{f_g}$. In the equilibrium, all h students are assigned to job g (actually, firms are indifferent between assigning any h student to job g and assigning to job m). Note that most of the qualitative results of the two-job model still hold in the three-job model: all the schools inflate to the extent such that their h student are of the same quality; more renowned schools inflate more than less renowned schools; all schools and firms are worse off in the equilibrium compared to the case without grade inflation.

However, the equilibrium outcome of the three-job model is different from that of the two-jobs in two aspects. First, in the three-job model all the schools inflate less than they do in the two-job model, since $\widehat{f_g} > \alpha$. The main reason is the existence of job m. Actuality, one can think of job m as a commitment device. Recall that in the equilibrium of the two-job model, all the schools are worse off compared to the case without grade inflation, since nobody gains at the expense of others. Collectively it is better for them to reduce the degree of inflation, but the problem is that individually each school has incentive to grade inflate more. The presence job m reduces each school's incentive to grade inflate, because if they grade inflate too much, their h students will be assigned to job m and getting a low wage due to the efficiency loss. But still individual school's incentive to grade inflate does not completely disappear: each school still wants to inflate as much as possible as long as their h student are assigned to job g. Here, firm's inability to wage discriminate within the job still has the bite.

Second, in the equilibrium of the three-job model $w_g > w_b$ since $\widehat{f_g} > \alpha$. But in the equilibrium model of two-job model $w_g = w_b$. In this feature the three-job model is more realistic than the two-job model, since in real word good there are wage premiums for good jobs.

We expect that the most of the qualitative results of our model still hold in a generalized n-job model, as long as there is a technology jump between the best job and the second-job. The unique equilibrium is that each school inflates to the extent such that firms are indifferent between assigning its h students to the best job or the second best job. In the equilibrium outcome, grade inflation still exists, more renowned school inflate more, and all the schools and firms are worse off compared to the case without grade inflation. But now the degree of grade inflation depends on the difference between the best job and the second best job. The bigger this difference, the bigger the equilibrium amount of grade inflation.

VI Discussion and Conclusion

Our goal was firstly to show that there is a very natural setting in which grade inflation could occur: adverse selection in the labor market. Secondly, we wanted to show that unlike price inflation, grade inflation is not "neutral", i.e., has "real" effects. With regard to the first goal, we argue that a free-rider problem is at the heart of grade inflation. Each university does not seek to maintain the collective reputation of the nation's cohort of college graduands. The university's ability to grade-inflate to milk this resource without regard for its consequences is the driving force of the model. This is disciplined only by firms' willingness to accept ones' h-grade students as truly high types. With regard to our second goal, we have shown that job mis-assignment results in some loss of social surplus. Welfare analysis demonstrates that all agents are better off in a world without grade inflation. We also show that the existence of disciplined schools will curb grade inflation by undisciplined schools to some extent, and reduce the welfare loss from mis-assignment, but will not eliminate it entirely. Our robustness analysis demonstrates that as long as there is a technology jump between the best job and the second-best job, grade inflation is still an equilibrium phenomena.

Our model delivers some other very stark implications which can be related to some interesting empirical regularities. We have shown that in equilibrium all schools will grade-inflate. Furthermore, grade inflation is contagious: others' inflationary practices will cause each school to do the same. We also try to relate grade inflation to (1) the reputation of universities, and (2) the extent of inequality in the skills requirement between good jobs and bad jobs. This is consistent with several empirical findings. The research cited in our introduction suggests that grade inflation is pervasive across higher education in America, and cuts across all institution types and across all fields of study (albeit to varying degrees). Furthermore, recent controversy in the popular press has centered around the suggestions that elite institutions have been more susceptible to grade inflation than others. These are consistent with our model which predicts that high reputation universities can afford to and will inflate more. What is also an interesting implication, in our opinion, and perhaps what lends a strong "economic" flavor to the issue of grade inflation, is that grade inflation might be related to rising inequality between skilled and unskilled workers; a trend we have observed in America in the past two decades. Some notable papers point out this trend. Lawrence Katz and Kevin Murphy (1992), in a famous paper, documented the rising skills premium in America from the 1960s to the early 1990s. The trend in grade inflation ever since the 60s appears to correspond to this time frame.

In our extended model, the introduction of an intermediate job can reduce the amount of equilibrium grade inflation. This may explains why firms sometimes use probationary periods for new employees. Probationary period for a good job might serve the function as an intermediate job: it increases firms' flexibility in job assignments and curbs grade inflation (elaborate more?)

To consider alternative theories of grade inflation, it is perhaps important to put grade inflation in its historical context. Education researchers have identified a "litany" of causes from which we describe four: the Vietnam war; affirmative action practices after the civil rights era; consumerist attitudes towards higher education; and faculty incentives.[17]

It is believed that higher education exemption during the Vietnam war formed the origins of grade inflation. It is argued that during that time, professors were loath to award low grades to male students lest they be drafted into the war. Because this period also coincided with the civil rights era, others have also argued that affirmative action may have contributed to grade inflation as universities admitted less-prepared students in the interests of diversity. Yet other reasons for grade inflation have been found in the changing attitudes towards higher education: especially the view that the university was increasingly commercialized. This led to the claim that students demanded high grades by virtue of the high tuition that they coughed up. Finally there are explanations rooted in the incentive effects of student evaluations on professors' tenure prospects and merit pay increases.

While these are set in a historical context, some of these explanations may provide fodder for interesting alternative stories to the one we tell, and to our knowledge, not all these phenomenon have been modelled. Actually, our model is related to the view of commercialization of universities. It is not a coincidence that grade inflation and the process of the commercialization of high education happened during the same period: from middle 60s to 80s. It was the change of universities objectives (commercialization) that unleashed the momentum of grade

[17] We refer the reader to Rosovsky and Hartleys' (2002) report for a concise discussion of these factors.

inflation.

One final point. One could argue that it is professors, not university administrators, who grade-inflate. However, universities are well aware of this trend but choose not to stop it. Actually, universities usually set some grading policies, like grade curves, which are followed by professors. Therefore we believe it is reasonable to treat university administrators and professors as the same agent. Our model highlights the "*external*" concerns of the university : its reputation, alumni relations, and labor market performance of graduands. We realize that a strategic interaction between professors and students may also generate interesting outcomes due to "*internal*" concerns. In particular, the effect of policies regarding tenure and merit pay increases on faculty performance may matter greatly; especially in settings where student evaluations comprise a significant component of faculty teaching performance. This is the subject of our next work.

References

[1] *Boston Globe*. "Harvard's Quiet Secret: Rampant Grade Inflation." October 7, 2001.

[2] *Boston Globe*. "Harvard's Honors Fall to the Merely Average." October 8, 2001.

[3] Chan, William, Hao Li and Wing Seun. "Why Are Grades So High," April 2003, working paper, the University of Hong Kong.

[4] Coate, Stephen, and Loury, Glenn C. "Will Affirmative-Action Policies Eliminate Negative Stereotypes?" *American Economic Review*, December 1993, 83,1220-1240.

[5] Juola, Arvo E. "Grade inflation in higher education. Is it over?" ED189129 (March 1979).

[6] Juola, Arvo. "Grade inflation in higher education:What can or should we do?" ED129917 (April 1976).

[7] Kuh, George and Hu, Shouping. "Unraveling the Complexity of the Increase in College Grades from the Mid-1980s to the Mid-1990s." *Educational Evaluation and Policy Analysis* (Fall) : 297-320. (1999)

[8] Levine, Arthur and Cureton, Jeanette S. *When Hope and Fear Collide: A Portrait of Todays College Student*. San Francisco:Jossey-Bass. (1999)

[9] Lizzeri, Alessandro. "Information revelation and certification intermediaries," *The Rand Journal of Economics*, Summer 1999; Vol. 30, Iss. 2; pg. 214-231.

[10] Ostrovsky, Michael and Michael Schwarz. "Equilibrium Information Disclosure: Grade Inflation and Unraveling," January, 2003, working paper, Harvard University.

[11] Rosovsky, Henry, and Hartley, Matthew. *Evaluation and the Academy: Are We Doing the Right Thing?*, American Academy of Arts and Sciences, Cambridge, MA. 2002.

[12] Sabot, Richard and Wakeman-Linn, John. "Grade Inflation and Course Choice." *Journal of Economic Perspectives*, Winter 1991, 159-170.

[13] *The Economist* "An Eye for an A". March 7, 2002.

[14] Tirole, Jean. "A Theory of Collective Reputations (with Applications to the Persistence of Corruption and to Firm Quality)," *Review of Economic Studies*, Vol. 63, No. 1. (Jan., 1996), pp. 1-22.

[15] Weiss, Andrew. "Job Queues and Layoffs in Labor Markets with Flexible Wage", *Journal of Political Economy*, vol. 88.3 (June 1980), 526-538.