Signal Reception and the Stagnation of Modern Education

Garrett M. Petersen[1]
Economics Department, Queen's University

July 30, 2013

[1]I would like to thank my advisor, Robin Boadway, for for his invaluable comments and guidance throughout this project. An earlier draft of this paper was presented at the *Symposium on Scholarship and a Free Society* at Chapman University on June 21, 2013. I would like to thank Art Carden and the others who made comments at that seminar. ©

Abstract

While the past century has brought rapid changes in almost all areas of life, high school curricula have been remarkably rigid. This paper argues that the problems of signal reception can explain excessive rigidity in the school curriculum. A model is developed to demonstrate how the incentives of university admissions departments to receive, or not to receive, certain educational signals imposes costs on schools and students who deviate from the standard curriculum. It is argued that this eliminates individual schools from the innovation process, leaving only large institutions capable of instituting curricular change, though these are not ideal innovators. Finally, possible solutions to the problem of curricular innovation are discussed.

Contents

1 Introduction

Perhaps nothing seems so clear to so many people as the shortcomings of modern education and the ways it ought to change. In the past decades, the world has changed, and dramatically, with expanding markets, widespread economic growth, and the steady march of technological progress creeping into almost every aspect of life. And yet, high school education has changed little in response.

This paper offers an explanation for the apparent rigidity in high school[1] curricula. The paper argues that this rigidity can be understood through the lens of signalling theory. Innovations can change the way things are taught, or they can change the actual content being taught, and I argue that there is a dearth of the latter type of innovations because new and different learning material sends an unrecognisable signal to the universities to which students apply. Finally, I relate this to the policy debates on programs to introduce more market competition into high school education.

The signalling model of education was developed by Spence (1973), who recognized that employers may not know a worker's marginal product when they hire him. Because of this, workers have incentive to signal their quality by doing things that only a good worker would do, such as spending a certain number of years in school, and that employers can observe and include in their hiring decisions. Lower-quality workers, in turn, have incentive to mimic

[1] Throughout this paper, "high school" is used to refer broadly to all secondary schooling institutions.

the signalling actions of higher-quality workers. Importantly, the signalling activities of workers must have differential costs depending on the workers' quality. Low quality workers face a higher cost to getting additional years of schooling than do high quality workers, so to separate themselves from low quality workers, high quality workers must get a level of education such that it would be too costly for low quality workers to gain by mimicking.

The Spence model implies that people get more than than the socially-optimal level of education. Education is seen, rather, as an unproductive arms race between workers. While signalling is not mutually exclusive with the notion that education improves workers' marginal product, the conclusion is that people over-invest in schooling from the perspective of economic efficiency. This inefficiency can be larger or smaller depending on the usefulness of the information gained in the process of signalling, and that's where innovation in the school curriculum becomes important. At worst, students could get an education without any value whatsoever, purely to broadcast to employers that they are capable of doing so.

This paper shifts focus away from the behaviour of the students sending educational signals, and rather focuses on the incentives of those who receive those signals. When there are costs associated with receiving a signal, the receiver, in this case a university admissions department, must get a benefit large enough to justify incurring the cost. This creates a peculiar sort of economy of scale in educational innovation: if a new curriculum is adopted on too small a scale, a lack of universities willing to interpret the grades issued under that curriculum imposes costs on the students. It is argued that this quashes innovation in curricula by excluding individual schools, public and private, from the innovation process.

2 Three Margins of Educational Innovation

There are three margins on which schools can innovate. The first is pedagogy, the second is the specificity of sorting, and the third is the quality of the curriculum. This paper is primarily concerned with the curriculum, so it is necessary to separate this margin from the other two.

Much has been written on the pedagogical quality of schools (eg. Ehrenberg and Brewer, 1994; Eide and Showalter, 1998; Hanushek and Rivkin, 2006). Whenever schools are compared on the basis of students' scores on standardized tests, the quality of each school's pedagogy, independent of the content being learned, is the variable under study. Innovations on this margin are productive to the extent that the content being studied is valuable. Otherwise, an improvement in teaching at one school or set of schools simply allows the students there to successfully mimic better students in the short run. In the long run, those to whom the students are signalling will adjust their beliefs.

Sorting can be thought of as a distinct service schools provide to their students. By testing students and distributing grades, educators allow each student to distinguish himself from his less skilful peers. More testing allows more specificity in sorting, and standardized testing allows students to distinguish themselves from a larger pool of peers. In increasing the specificity of sorting, the less skilful are made worse off, so innovations on this margin may or may not be productive[2].

The third margin on which schools can innovate is the curriculum. Schools can improve their curriculum by making the content more applicable to the work their students will eventually perform, offering improvements to their human capital, or by simply making the content more enriching and enjoyable in itself. Innovations on this margin are generally productive, so the apparent fact that they are slow in coming is troubling and requires special explanation.

3 Lamentations on Math Education

There is no one thing called human capital. The human capital developed in high school is a disparate bundle of skills and information, some of which may have a large effect on students' productivity while some may not. If schools responded efficiently to changes in the larger economy, we would expect them to teach the bundle of skills for which the marginal benefits are at least as great as the marginal costs associated with learning them.

For instance, if there were a learned skill that had a close substitute,

[2]Stiglitz (1975) argued that sorting can be efficient since it allows a better match between workers and jobs. In the context of this paper, where students are signalling to universities, a better match between students and universities could produce social gains. Indeed, in the model in section 5, the ability of universities to reject the worst students does produce efficiency gains.

and the price of that substitute fell precipitously, we would expect schools to shift away from teaching that skill. There is just such a skill, and that is the ability to quickly perform algebra and arithmetic by hand. This has a close substitute in machine computation, which is dramatically cheaper than it once was. In an age when most grocery stores use machines to count back exact change to customers, shouldn't schools spend less time teaching students to compute differences in their heads?

As computers have reduced the value of the ability to do arithmetic by hand, they have increased the value of other forms of math, those for which computing power is complimentary. Even a basic statistical regression would be prohibitively time-consuming for a human to perform by hand, and yet with a cheap laptop that same human can run dozens of regressions in seconds, and use them to answer questions about the real world. Thus there have been changes in the relative values of different sorts of math skills.

These changes have not gone unnoticed. There is a movement spearheaded by Conrad Wolfram (see his TED talk, Wolfram, 2010) to change schools over to a program of computer-based math. Wolfram views math as a four-step process: (1) posing the right question, (2) formulating that question in math, (3) computing the answer, and (4) verifying that answer. He claims, very plausibly, that 80% of students' time is spent on (3), while (1), (2), and (4) are neglected. The computing done in step (3) is done vastly better by computers, so Wolfram would like to see schools shift to teaching the mathematical way of thinking rather than focusing on computation by hand (Wolfram, 2010).

Wolfram recognizes the difficulty in changing the math curriculum. He does not view the change from the current math curriculum as something that can be achieved through incremental reform, rather he views it as a "chasm" to be leapt across.

Paul Lockhart is even less generous to the standard math curriculum. He argues persuasively in his unpublished essay, "A Mathematician's Lament" (2002), that what goes under the name of "math" in schools is not something any mathematician, living or dead, would recognize as math. School "math" is to math as memorizing and labelling paint swatches is to painting[3].

> "The most striking thing about this so-called mathematics curriculum is its rigidity. This is especially true in the later grades. From school to school, city to city, and state to state, the same exact things are being said and done in the same exact way and in the same exact order. Far from being disturbed and upset by this Orwellian state of affairs, most people have simply accepted this"standard model" math curriculum as being synonymous with math itself." (Lockhart, 2002, pg. 14)

Both Lockhart and Wolfram recognize the rigidity of the standard model math curriculum. The observation that the high school curriculum is particularly rigid across time is not new, however. In 1990, Smith and O'Day complained that the processes and content of instruction were "little different from what they were in 1980 or in 1970" (pg. 233).

[3]It should be noted that Lockhart's view, that math education as it stands has little to no value, is among the more extreme views. However, some level of dissatisfaction with education as it now stands is commonplace. Although the status quo undoubtedly has defenders, this author is not aware of any defences that have been made explicitly in print.

Although many people, including those cited here, have observed the lack of change in school curricula and the shortcomings of the curriculum as it now stands, there has been little discussion of the systematic reasons why we might expect this rigidity more in schools than in other areas of life. Of particular interest to economists is why the invisible hand of the market has not adapted mathematics education to students' needs without the need for reform movements like that headed by Wolfram. What Lockhart hints at, although he does not use these terms, is an inefficient signalling equilibrium in which university admissions departments recognize what good grades in the "standard model" math curriculum signal about a person, and any who deviate bear the cost of sending signals that nobody recognizes[4].

Math could be more fun, more engaging, more relevant, and less soul-crushing, but it doesn't become those things because the first to make the switch would bear too great a cost. An A+ in AP calculus, and in the sequence of courses leading up to it, is a sure sign to any employer or university that a student can work hard and follow instructions even when the task is as boring and pointless as memorizing reams of notations and formulas and applying them without the motivation of solving anything of any interest, importance, or beauty. That's a powerful signal, and failing to send it is a good way to lose out on a big scholarship from your university of choice.

This is not to say that universities are being unreasonable. Their incentive

[4] "Nothing looks better than Advanced Paint-by-Numbers on a high school transcript," is spoken by an imaginary art teacher, Lockhart's colourful way of saying that math, as taught, is a useless task justified purely for signalling purposes (Lockhart, 2002, pg. 2).

is to separate good students from bad ones, and the costs of interpreting any signal may be too high to justify doing so given that few students are sending that signal.

4 Explaining Curricular Rigidity

Before launching into the signal-reception model of curricular rigidity, it is worthwhile to discuss the validity of other potential explanations.

One possible explanation is that the high cost of retraining teachers to teach a new curriculum prevents a change from occurring. One can imagine a scenario when the cost of retraining a teacher is too high to justify the benefit of doing so. However, this cannot explain why new cohorts of teachers continue to be trained to teach the old curriculum. It does, however, explain why incumbent teachers have little incentive to push for innovation. Each new teacher trained to teach a more desirable curriculum devalues the skills of those who are heavily invested in expertise in teaching the old curriculum. Thus the influence of teachers' unions, organizations that exist to protect the interests of incumbent teachers, is distinctly anti-innovation. If we were to attribute the rigidity in curricula to teachers' unions, we should expect innovation to continue in private schools. However, private schools differ from other schools mainly in their pedagogy; private school curricula tend to differ little from those of other schools with respect to the core concepts under study.

It's the ubiquity of the standard curriculum that's so puzzling. If teachers' unions or bureaucratic inertia are blamed, one can point to a school somewhere that adheres to the standard curriculum despite being exempt from union rules and from the management of government bureaucrats. The rigid adherence to standard curricula in independent, privately-run schools requires further explanation, and that is what the signal-reception model offers.

5 A Signal-Reception Model

The signal-reception model describes a world where high schools choose a curriculum, and students use the grades assigned to them under that curriculum to signal their quality to universities in different cities.

Students are differentiated by their quality, θ, and by the city in which they initially live, of which there are two. Student quality in each city is continuously distributed according to $f(\theta)$ with a range from 0 to 1, and the student population of each city is normalized to 1. There is one university in each city, and these universities always know students' city of origin, but they do not necessarily know each student's quality. As I seek to demonstrate the difficulties associated with those on the receiving end of educational signals, I assume that it costs a student nothing to signal his quality, and it would cost him an infinite amount to mimic a different student's signal, effectively imposing a separating equilibrium by assumption.

Students value university admissions independently of their θ; they value admission to the university in their city at V_{home} and they value admission to the university in the other city at V_{away}, where $V_{home} > V_{away} > 0$[5]. Students' reservation wage is 0. Universities get a benefit, $B(\theta)$, from admitting a student of a given quality. $B(\theta)$ is strictly increasing in student quality, such that $B(0) < -V_{home} < 0 < B(1)$, and that

$$\int_0^1 [B(\theta) + V_{away}] f(\theta) d\theta > 0. \tag{1}$$

These assumptions ensure, respectively, that some students are accepted and others are rejected when universities know students' quality, and that all students are accepted when neither university knows students' quality.

Universities act to maximize their profit, π, which is the sum of the benefits the university gets from students and the tuitions paid by those students.

[5] That all students get an equal benefit from admittance to university is not an entirely realistic assumption. In reality, we should expect better students to benefit more from university than worse students do. However, if there were a strict monotonic relationship between student quality and the benefits from university admittance, then universities could elicit student quality through the price mechanism, eliminating the need for admissions departments. Since admissions departments do exist, we can infer that the relationship between student quality and student benefits from university are not so perfect that prices can be the only allocation mechanism. The most realistic assumption would be to have benefits generally increasing in student quality, but with some random error so that the price mechanism would imperfectly sort students. However, this would add a great deal of mathematical complication and would deliver much the same result as the constant value assumption, so I use the latter.

5.1 Competition Between Universities

I proceed by first solving for the tuition rates, net of scholarships, in the market for university admissions when neither university knows students' quality, when only one university knows students' quality, and when both universities know students' quality. I will then use these values to determine the value universities get from knowing students' quality, and the value students get from being able to signal their quality.

5.1.1 Neither University Knows Student Quality

I begin by solving for the case where neither university knows students' quality. Competition for students is Bertrand. This means that each university attracts the students from its home city by charging the highest price such that the other university cannot profitably undercut. The lowest price a university from city j could charge to students from city i is

$$T_j = -\int_0^1 B(\theta)f(\theta)d\theta$$

as that would yield a profit of $\pi_j = 0$ from these students. Therefore, the highest price university i can charge to students from city i without being undercut is

$$T_i = V_{home} - V_{away} - \int_0^1 B(\theta)f(\theta)d\theta.$$

This is the price both universities charge to students from their respective cities, each earning a profit of $\pi = V_{home} - V_{away}$.

5.1.2 One University Knows Student Quality

In the case where one university knows students' quality, there is a dual problem in admitting a student of given quality. If the university with knowledge of students' quality admits a student of quality θ', the university must consider both the direct costs and benefits, given by $T_i(\theta') + B(\theta')$, and the effect on the maximum price the competing university can charge, as this affects the price the university with knowledge can charge to other students.

Suppose university i knows students' ability in city i, but university j does not. Then university i maximizes its profits by price discriminating between students of different quality such that university j does not enter the market. It is reasonable to assume that the tuition function $T_i(\theta)$ is decreasing in θ, as high-quality students are more valuable and both universities are willing to pay more for them. Thus university i charges a student of value θ' the highest price such that university j will not enter the market and capture all students with $\theta \leq \theta'$, and that the student will still choose to attend university. The minimum value of T_j university j could charge all students of quality $\theta \leq \theta'$ is $T_j = -\int_0^{\theta'} B(\theta) f(\theta) d\theta$. Thus a student of quality θ' is charged tuition

$$T_i(\theta') = \max\{V_{home} - V_{away} - \int_0^{\theta'} B(\theta) f(\theta) d\theta, V_{home}\}$$

by university i so long as $T_i(\theta') \geq -B(\theta')$, otherwise the student is simply rejected.

If university i knows the quality of students in city j, but university j does not, university i can use price discrimination to capture students in city j. However, in this case university i must earn a loss on some students in order to drive down the minimum tuition university j is willing to charge.

I once again assume that $T_i(\theta)$ is decreasing in θ. The highest-quality students have the largest positive effect on the expected benefit to university j of charging a given tuition, so if it is worthwhile to take a loss on some students to earn on others, the highest-quality students must be among those competed away, so the set of students admitted by university i must be bounded above by 1.

Thus I solve for the tuition function in the same way I did previously. University i charges a student of quality θ' the most it can such that university j does not enter at a price high enough to compete away all students of quality $\theta \leq \theta'$. This means the tuition function is

$$T_i(\theta') = \max\{V_{away} - V_{home} - \int_0^{\theta'} B(\theta) f(\theta) d\theta, V_{away}\}$$

unless θ' is sufficiently low that $T_i(\theta) + B(\theta) \leq 0 \; \forall \; \theta \leq \theta'$, in which case the student is rejected. If this tuition scheme earns a positive profit, university i will enter the market.

University j may still find it profitable to admit the students rejected by university i. Charging a price of V_{home}, university j can earn a profit of $\int_0^{\underline{\theta}} [B(\theta)+V_{home}]f(\theta)d\theta$ where $\underline{\theta}$ is the lower-bound for admission to university i. If this profit is negative, university j will choose not to enter the market.

If university i does not find it profitable to enter the market, university j will admit all students at a tuition rate of

$$T_j = V_{home} - V_{away} - \int_0^1 B(\theta)f(\theta)d\theta$$

as in the case where neither university knows students' types.

5.1.3 Both Universities Know Student Quality

In the case where both university i and university j know the quality of all students, the universities must compete for each student separately. This means that for a student of quality θ', the university in his home city will charge him the highest tuition such that the other university cannot profitably undercut. The other university cannot charge less than $T_j(\theta) = -B(\theta)$, as at this tuition its profits are 0. This means that the home university will charge $T_i(\theta) = \max\{V_{home} - V_{away} - B(\theta), V_{home}\}$ so long as $T_i(\theta) \geq -B(\theta)$, otherwise the student is simply rejected. The universities face a parallel problem with respect to students in city j,

5.2 Schools' Signalling Problem

High schools must consider universities' incentives when choosing a curriculum for their students. I assume that universities face a one-time learning cost, L, to find out what the grades issued under a given curriculum mean. In other words the university can pay L once for a given curriculum, and from then on it knows the quality, θ, of any student that has studied under that curriculum.

We would expect an improvement in a school curriculum to either deliver a benefit directly to the students, or to change universities' benefit from admitting the students (ie. changing the shape of $B(\theta)$). However, if $B(\theta)$ were to increase, the competition between universities would lead the benefits to be passed on to the students in the form of lower tuitions. Thus improvements that increase $B(\theta)$ and improvements that deliver value directly to students are equivalent in all respects except that increases in $B(\theta)$ make a larger proportion of students enter university in the cases where universities are aware of student quality. This complicates the analysis without any qualitative changes to the results, so I consider a curricular change that delivers benefits directly to students.

Suppose that all students in both cities study under one curriculum recognized by both universities, and suppose that a school in one city is considering changing to a new curriculum. Under the status quo, students' quality is known to both universities. The school contains n students, and these n students are a representative sample of the entire student population (ie.

their quality is distributed according to $nf(\theta)$).

If the new curriculum is adopted, the universities have two choices: they can learn to interpret the signals sent under that curriculum, or they can choose not to.

		Other University	
		Learn	Don't Learn
Home University	Learn	$\pi_h^{L,L} - L, -L$	$\pi_h^{L,DL} - L, 0$
	Don't Learn	$\pi_h^{DL,L}, \pi_o^{DL,L} - L$	$\pi_h^{DL,DL}, 0$

Table 1: Payoff table for universities choosing whether or not to learn to interpret the new grades when a school in one city changes its curriculum. Subscripts distinguish universities (h for home and o for other) and superscripts distinguish cases (eg. L, DL indicates home university has learned while the other university hasn't). $\pi_h^{L,L}$, $\pi_h^{L,DL}$, and $\pi_h^{DL,DL}$ are strictly positive, while $\pi_h^{DL,L}$ and $\pi_o^{DL,L}$ are merely non-negative as either university may choose not to enter in this case. This eliminates the possibility that $\{Learn, Learn\}$ is a Nash equilibrium.

Since the other university gets no benefit for learning if the home university learns, $\{Learn, Learn\}$ cannot be a Nash equilibrium, though any of the other three outcomes could potentially be. The graph in figure 1 shows students' utility depending on universities' choice to learn or not learn, when the new curriculum offers no additional benefit. Universities' potential profits from students in the program are proportional to n. If there are enough students in the program to entice one university to learn, no student is made better off, and many are made worse off by the change.

If there are so few students in the new program that neither university has incentive to learn, then better students are made worse off under the

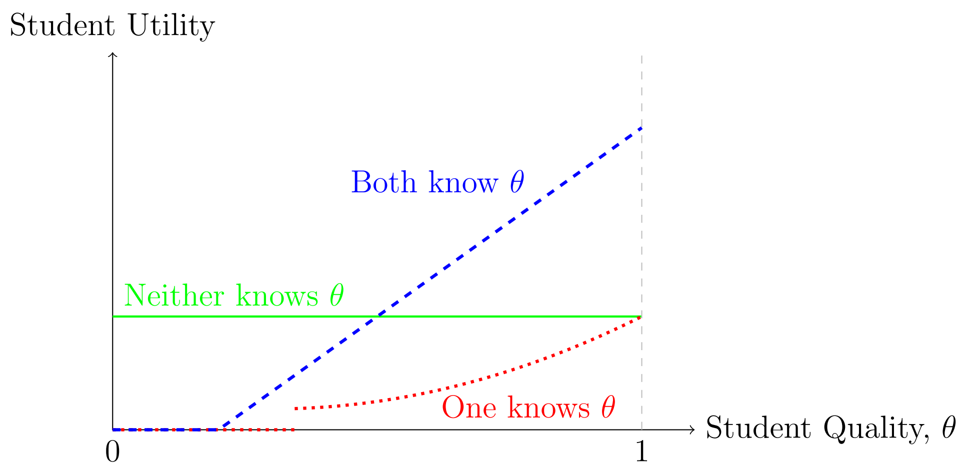

Figure 1: Student utility levels when $f(\theta) = U[0,1]$, $B(\theta) = \theta - 0.3$, $V_{home} = 0.2$, $V_{away} = 0.1$, and the new curriculum offers no additional benefit to students. The case where both universities know θ is the utility students get when both universities can interpret their grades, as is the case under the status quo curriculum, while the other two lines represent the utility students get when one or both universities do not know θ, as will be the case if they are changed to a new curriculum in only one city.

new curriculum, while worse students are made better off. However, this leads to an adverse-selection problem. If students can change schools in the long run, the best students will leave the school for schools still teaching the old curriculum, driving down the average quality of students in the school teaching the new curriculum. When the average quality of students in the school falls, the university charges higher tuition to the remaining students, leading more to exit the school, until eventually all benefits are erased.

Thus if any number of students in a given city are changed to a new curriculum, they will be made worse-off in the market for university admissions. Even if there are offsetting benefits associated with the new curriculum, these benefits must be sufficiently large to compensate students for this loss, or the new curriculum cannot improve students' welfare.

Figure 2 demonstrates the minimum additional benefit to students the new curriculum must have to avoid the adverse-selection problem described above. Since the adverse selection problem begins with the exit of the best student, the additional benefit of the new curriculum must be sufficiently large to compensate the best student for the loss of his scholarship. Otherwise students leave until only those with very low benefits from signalling, those of the lowest quality, remain. This model implies, therefore, that remedial programs do not face the same rigidity as do programs aimed at the general student population, since low quality students have less to lose from failing to signal.

Figure 3 demonstrates the case where one university learns the new cur-

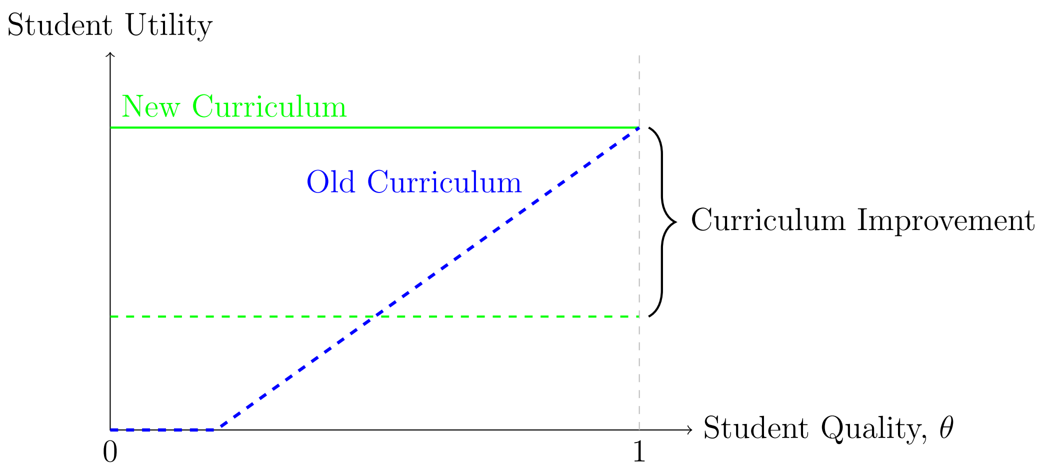

Figure 2: In the case where neither university is expected to learn to interpret the new curriculum, the additional benefit to students from the new curriculum must be high enough to compensate the best student for the loss of his scholarship to avoid the adverse-selection problem. If the new curriculum offers a benefit to students of any less than the amount shown, the best students will self-select out of the school with the new curriculum, driving down the average quality of students and increasing the tuition universities charge to the students who remain.

riculum. The best students once again self-select out of the school that introduces the new curriculum unless the benefits are exceptionally large. Once again, the model predicts less curricular rigidity in remedial programs than in programs aimed at the general student population. If the best students are anticipated to self-select out of the program, the university is less likely to find it worthwhile to learn the new curriculum, making the no-learning equilibrium the more likely scenario.

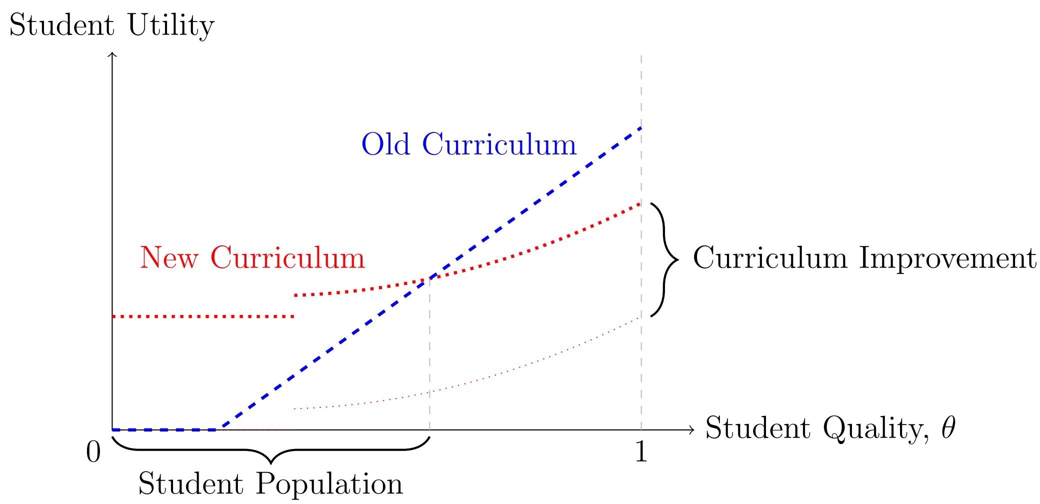

Figure 3: In the equilibrium where one university learns, the additional benefit of the new curriculum must be very large to be worthwhile for the best students. If, as is shown, the new curriculum only delivers a small benefit to students, the best students will self-select out so they can have the benefit of having both universities aware of and competing for their quality.

The only way to change to a new curriculum without encountering the adverse-selection problem of the no-learning equilibrium or the price discrimination problem of the equilibrium in which one university learns is for a sufficient number of students of quality in *both* cities to change to a new curriculum such that universities are motivated to learn to interpret the signals sent under the new curriculum. Neither university will consider its effects on competition in the other city, as its profits are zero there in any case, but it will learn if the profits within its own city are enough to justify doing so.

6 Innovation by Decree?

The model leads to a stark conclusion: improvements in high school curricula must be made on so large a scale that it pays multiple competing universities to learn to interpret the new curriculum, or else the benefits from an improved curriculum are likely to be swamped by the costs students face in the market for university admissions. This effectively excludes individual private or public schools, or even government bodies in charge of public schools in smaller regions, from the innovation process.

Without many people experimenting and innovating on a small scale, schooling cannot be as innovative as it would ideally be. A bottom-up, decentralized process wherein many ideas are tried and refined or rejected serves to sift out all but the best ideas; it is a process like this that produces the constant stream of technological advances that define the modern world. Sadly, with individual schools removed from the potential pool of innovators, this process cannot occur.

The only potential innovators left standing are large institutions with control over many schools in many regions. In other words, if there are to be curricular changes in high schools, they must stem either from central governments or from teachers' unions. Teachers' unions' lack of incentive to innovate was discussed in section 4, so I will focus here on politicians' lack of incentive to enact positive reform.

Elected politicians are motivated to respond to voters' preferences over school curricula, but there's little reason to think that these preferences accurately reflect which curriculum is best. Voters are rationally irrational (Caplan, 2001) about the best policies, as no individual vote is likely to impact the policies enacted. Voters can indulge their biases, with little to no cost to themselves, when casting their ballots. Insofar as voters have a bias towards the status quo, politicians will keep the school curriculum the way it is, and if voters' biases favour something other that the status quo there is no reason to think that it will be an improvement.

A politician would rationally enact a curricular change that went against his voters' beliefs if he thought there would be significant and undeniable benefits within his term in office, so voters would change their minds in his favour. However, the positive effects of an improved curriculum are unlikely to be felt until after the students have entered the workforce, too late for a politician looking only as far forward as the next election.

This line of argument paints teachers and politicians as narrowly self-interested. However, it is possible that elected politicians and the teachers in charge of teachers' unions are not coldly self-interested, and that they genuinely want to improve schooling for the students' benefit. This leads to their second difficulty as innovators: the knowledge problem.

One might think that a teachers' union would have the expertise to choose the best curriculum. Teachers are experts in one thing: teaching. However, expertise in teaching is not sufficient to choose the ideal school curriculum. To do so, one should also be an expert on all the tasks students could eventually use their knowledge to perform, so the marginal cost of gaining that knowledge could be equated with the expected marginal benefit. This is not only beyond teachers' expertise, it is beyond the expertise of any human.

As Hayek put it, "the knowledge of the circumstances of which we must make use never exists in concentrated or integrated form but solely as the dispersed bits of incomplete and frequently contradictory knowledge which all the separate individuals possess" (Hayek, 1945). The problem of choosing the right curriculum to learn in school is a problem of integrating the knowledge of the value of goods and services, held by the people who consume them; the knowledge of the marginal usefulness of the various skills one can learn in school in the production of those goods and services, held by the managers of firms; the knowledge of the costs associated with learning, held by teachers and school administrators; and the knowledge of students' individual talents and inclinations, held by the students themselves.

The knowledge of the value of goods and services and of the marginal product of various skills are transmitted, albeit imperfectly, through the wages paid to workers in the market. Students, knowing their own interests and having access to information on the wages paid to different sorts of workers, can narrow down their career paths and figure out the best sort of schooling for them. However, neither teachers' unions nor central governments are in a position to elicit that information from students, as they do not compete in a market against alternative schools with alternative curricula for which students could express their preference.

7 Can We Have Market-Driven Innovation?

I have painted a bleak picture of the possibility of positive reform in school curricula. Indeed, it seems that as things currently stand, schools are unlikely to see much improvement.

If all parties—schools, universities, and students (including those yet unborn)—could contract freely without transaction costs, they could reach the best outcome (Coase, 1960). Clearly such a contract is impossible. However, if private schools could contract with universities, reimbursing their learning costs to allow a change to an improved curriculum, or if enough private schools could contract with each other to incentivize universities to bear the learning costs themselves, those private schools could increase their desirability and thus their profits until other schools succeed in copying the new curriculum. This would allow some market-driven innovation to occur, though it would not be the economically efficient level of innovation because of the eventual free riding of other schools.

Thus market-driven innovation can occur if a large enough number of private schools can contract so as to entice universities to learn to interpret new curricula. Hypothetically, a market could arise to fill this need. Changing to a new curriculum would require the (re)training of teachers, the writing of new textbooks, and the administering of new standardized tests. If firms could arise to provide these services, in addition to acting as intermediaries between schools and universities, they could compete to offer the curriculum most desired by students. The appearance of such a market would solve the rigid curriculum problem.

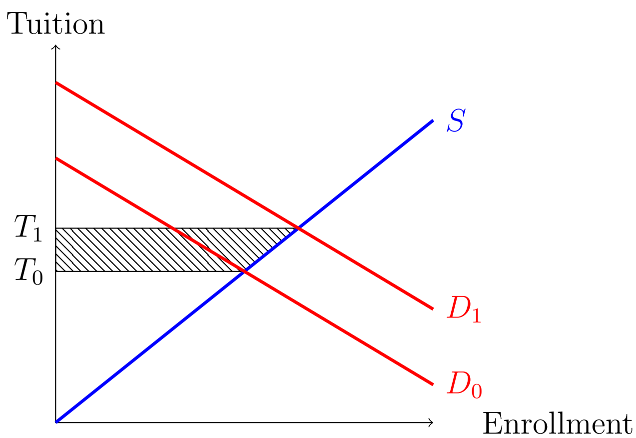

Figure 4: If a new curriculum increases students' demand for entrance into a school, the shaded area represents the maximum the school will pay to change to the new curriculum.

It seems that such a market has not arisen because the profits to be had are not sufficient to justify the entry cost. Figure 4 shows the maximum a private school would pay to improve its curriculum. This minus the costs of doing so is the maximum profit a firm could earn if it were to sell a new curriculum to the school. The entry cost is that of developing the new curriculum and of contracting with universities to ensure they recognize the new curriculum in their admissions process. Thus, if the firm is to enter the market, there must be a sufficient number of schools willing to be its customers. Due to the public schools' lack of incentive to innovate, described in section 6, the demand for innovative new curricula must come from private schools, and there simply aren't enough of them.

8 Conclusion

There is an ongoing policy debate about introducing more markets into education (see for instance Neal, 2002; Ladd, 2002), whether through school vouchers or other such policies, and the line of reasoning presented in this paper has a somewhat surprising implication for that issue. Because of the curricular rigidity herein described, private schools compete on the pedagogical margin but not on the curricular margin. However, if there were enough schools competing in a market to allow the creation of a market for curricula, such as that described in section 7, there would be competition and innovation on both margins.

This implies a non-linear benefit to markets in education. A small expansion in the amount of education delivered through markets rather than government is likely to create a smaller proportional benefit that a much larger expansion, if the large expansion could be sufficient to allow competition between curricula.

The central observation in all this is that we should expect a school system that is predominated by public schools to be no more innovative in developing new and better curricula than any industry run entirely by government. The market for private schooling, such as it is, is too small to be the source of new ideas and change. Reformers must aim to change the minds of politicians or the public to improve the school curriculum through the political process since, as things stand, the curricular rigidity problem prevents the market from delivering these improvements.

References

Bryan Caplan. Rational irrationality and the microfoundations of political failure. *Public Choice*, 107(3-4):311–331, 2001.

R. H. Coase. The problem of social cost. *The Journal of Law and Economics*, 3:1–44, October 1960.

Ronald G. Ehrenberg and Dominic J. Brewer. Do school and teacher characteristics matter? evidence from high school and beyond. *Economics of Education Review*, 13(1):1–17, 1994.

Eric Eide and Mark H. Showalter. The effect of school quality on student performance: A quantile regression approach. *Economics Letters*, 58:345–350, 1998.

Eric A. Hanushek and Steven G. Rivkin. Teacher quality. *Handbook of the Economics of Education*, 2:1051–1078, 2006.

F. A. Hayek. The use of knowledge in society. *The American Economic Review*, 35(4):519–530, September 1945.

Helen F. Ladd. School vouchers: A critical view. *The Journal of Economic Perspectives*, 16(4):3–24, Fall 2002.

Paul Lockhart. A mathematician's lament, 2002. URL `http://cct-dev.edc.org/userfiles/files/LockhartsLament.pdf`.

Derek Neal. How vouchers could change the market for education. *The Journal of Economic Perspectives*, 16(4):25–44, Fall 2002.

Marshall S. Smith and Jennifer O'Day. Systemic school reform. *Journal of Education Policy*, 5(5), 1990.

Micheal Spence. Job market signaling. *The Quarterly Journal of Economics*, 87(3):355–374, August 1973.

Joseph Stiglitz. The theory of "screening," education, and the distribution of income. *The American Economic Review*, 65(3):283–300, June 1975.

Conrad Wolfram. Stop teaching calculating, start teaching math, October 2010. URL `http://www.computerbasedmath.org/resources/reforming-math-curriculum-with-computers.html`. TED talk, transcript retrieved on June 22, 2013.