DRAFT (3/24/06): *Please do not cite without permission*

Academic Freedom, Intellectual Bias, and "the Truth": What Can Bayesian Statistics Teach Us?

Mark Hopkins
Department of Economics
Gettysburg College
Gettysburg, PA 17325
mhopkins@gettysburg.edu

"Every individual, whether conservative or liberal, has a perspective and therefore a bias. Professors have every right to interpret the subjects they teach according to their individual points of view. That is the essence of academic freedom. But they also have professional obligations as teachers, whose purpose is the instruction and education of students, not to impose their biases on their students as though they were scientific facts."

- David Horowitz

I. Introduction

Several state legislatures are now considering versions of a student's "bill of rights" that would regulate academic activities at state-funded universities. These rights, as enumerated by one group supporting this legislation, Students for Academic Freedom (SAF), include the following (SAF, 2006):[1]

(a) "Students will be graded solely on the basis of their reasoned answers…not on the basis of their political and religious beliefs."

(b) "Curricula and reading lists in the humanities and social sciences should reflect the uncertainty and unsettled character of all human knowledge in these areas by providing students with dissenting sources and viewpoints where appropriate….Academic disciplines should welcome a diversity of approaches to unsettled questions."

(c) "Exposing students to the spectrum of significant scholarly viewpoints on the subjects examined in their courses is a major responsibility of faculty. Faculty will not use

[1] Students for Academic Freedom actually use the term "Academic Bill of Rights" (SAF, 2006) and call for protections of "professors, researchers and students" to "teach and to learn." In practice, however, institutional protections already exist to guarantee academic freedom and protect speech for faculty, including tenure, outside peer review, grievance procedures, and non-discrimination laws, and there is little or no evidence that these do not function effectively. In this paper, therefore, I will narrow my focus to the primary distinguishing characteristic of the legislation: that students should have a "right" to "freedom" in "learning." In particular, of the eight "rights" listed by SAF, the first two – pertaining to faculty protections already covered by existing non-discrimination law – were omitted for the sake of brevity and focus.

their courses for the purpose of political, ideological, religious or anti-religious indoctrination."

(d) "Selection of speakers, allocation of funds for speakers programs and other student activities will observe the principles of academic freedom and promote intellectual pluralism."

(e) "An environment conducive to the civil exchange of ideas being an essential component of a free university, the obstruction of invited campus speakers, destruction of campus literature or other effort to obstruct this exchange will not be tolerated."

(f) "….Academic institutions and professional societies should maintain a posture of organizational neutrality with respect to the substantive disagreements that divide researchers on questions within, or outside, their fields of inquiry."

This paper will take as its motivation the principles elaborated in (b) through (d). Since it would be very difficult to find anyone who would to disagree with (a) in theory, and grievance procedures and protections to address this problem exist in practice at almost every academic institution, it will not be considered further. This is true for (e) as well, although admittedly the activities of students are generally much harder to regulate. Similarly, the principle outlined in (f) goes to the mission and organization of academic institutions themselves, and should therefore, I believe, be treated separately from the other principles. The rationale for this is that an argument can be made that students, having committed themselves to a particular academic institution, should have some guarantee of intellectual fairness. In the transparent and highly competitive market for higher education, however, the principle of *caveat emptor* most certainly applies: if a student admitted to Harvard views the faculty there as being "too liberal" then they are free not to attend. Any obviously discriminatory institutions, like Gary Becker's (1971) discriminatory firms, are apt to lose market share and eventually be driven out of the market. Nevertheless, the conclusion to this paper will revisit this point, and offer some discussion of what imperfections may exist, if any, in the market for higher education, and whether survey data that suggest that the number of academics that are registered Democrats is higher than the national average may reflect barriers to entry or simply selection bias in career choice.

The justification stated by SAF for the rights outlined in (b) through (e) are two-fold: "securing the intellectual independence of faculty and students" and "protection of intellectual diversity." The first principle is a bit unclear since it is not stated from whom this independence is to be granted: presumably it is not independence from politicians, since the point of the legislation being proposed is to grant the elected government oversight of the educational process. Moreover, are faculty and students supposed to be intellectually free of each other? This raises many operational complications since the activity at hand involves cooperative learning processes. Similarly, the lack of specificity regarding the principle of "intellectual diversity" is problematic, since it reflects an outcome rather than a process. How is this diversity to be achieved? If intellectual diversity implies nothing more than the absence of dogmatic constraints or political influence on thought and speech, then it is manifestly sensible. If such a "diversity requirement" requires heavy-handed regulatory influence over the marketplace for ideas, however, then its value is more in question. Understanding the meaning and implications of this principle will be central to the analysis in this paper.

In short, discussing the idea of a students' bill of rights in broad terms is inherently difficult, since the principles and rationale for some are redundant and/or uncontroversial while with many others "the devil is in the details." It seems clear however that the intrinsic controversy rests on the meaning and significance of several key ideas:

- Human knowledge, even scientific knowledge, is much better described as a set of "beliefs" than as "truth."

- A *belief* differs from the *truth* in three key respects: first, a truth is a certainty, while a belief acknowledges uncertainty; second, a truth is constant, while a belief is mutable; third , truths are objective while beliefs are subjective.

- Although beliefs can be shaped by data, two individuals rarely observe exactly the same lifetime of data and, therefore, rarely have exactly the same beliefs. If direct observation of the same data is not possible, it is not immediately clear how one person can (or should) convey their knowledge (beliefs) to another.

This paper will make an attempt to explore these issues by posing six specific sets of questions and then attempting to answer them within a formal statistical framework. These questions are as follows:

Q1. What is the meaning of "truth" in a context of uncertainty and to what extent can it be viewed objectively? Can evidence be presented in an ideologically-neutral manner?

Q2. What is the exact meaning of "bias" in the context of an uncertain truth?

Q3. What is the of "ideology" in the classroom and in the educational process? What are our pedagogical goals, and what should they be? In particular, should or should not professors be trying to convey "the truth" as they see it?

Q4. What problems do asymmetric information create in the classroom? Can the perception of "liberal bias" among students lead both to misperceptions of the truth and/or the emergence of an (actual) "conservative bias" in the classroom?

Q5. What conflict exists between the conventional definition of academic freedom as the right of professors to express their beliefs freely and the newer definitions that also include the right of students to express their beliefs freely? Do trade-offs exist from a welfare standpoint?

Q6. Is there reason to believe the marketplace for ideas is not efficient? Do structural and institutional barriers to entry exist that create a market failure?

II. Bayesian Statistics as an Analytical Framework

Questions that have been raised regarding *academic freedom* span a wide array of topics, from the scope of what is covered to the manner in which it is taught. A single paper cannot hope to address all of these questions satisfactorily, and so the analysis offered in this paper will address a single aspect of the debate: how to think about and handle the presentation of contested hypotheses. Instances often arise in which uncertainty exists regarding which of two (or more) hypotheses is "correct." I use quotations marks in this case because we have not yet established what is meant by a correct hypothesis: the hypothesis that is more consistent with the data, or the hypothesis that is true? The second definition is clearly more satisfying, but only the first is something that, given some underlying assumptions, can be reasonably and objectively stated.

These two concepts of what it means for a hypothesis to be *correct* are distinct, but not unrelated. This paper argues that Bayesian statistical theory can help us to understand their relationship, and thus to understand the decision problem that faces educators, students, and policy makers today. For those unfamiliar with Bayesian statistics, the following section offers a cursory introduction to the concepts. Those already familiar may wish to skip directly to Section III.

A. *Bayes Rule*

The Bayesian statistical tradition derives its name from Bayes Rule, which states that the probability of an event A, conditional on knowledge that event B has already occurred can be expressed as follows:

(1)
$$\Pr(A \mid B) = \frac{\Pr(B \mid A)\Pr(A)}{\Pr(B)}$$

The intuition behind this is quite simple, and can be seen using a Venn Diagram as in Figure 1. Imagine that every point inside the unit square has an equal probability of occurring (call it a possible "state of the world").[2] The probability that event *A* occurs can thus be represented by the total area inside the circle labeled "*A*" and the probability that event *B* occurs can be represented by the area in the circle labeled "*B*." The area inside both circles is the *intersection*, $A \cap B$. Clearly if we know *B* has happened, then two things are true:

1) if *A* also occurs, it must be that the state of the world lies within the intersection, $A \cap B$; and

2) the only possible states of the world that we can now be in lie in *B,* not *U.*

[2] Note that, as a 1x1 square, the total area of the square is equal to 1. That is, the square represents the universe of all possible states of the world.

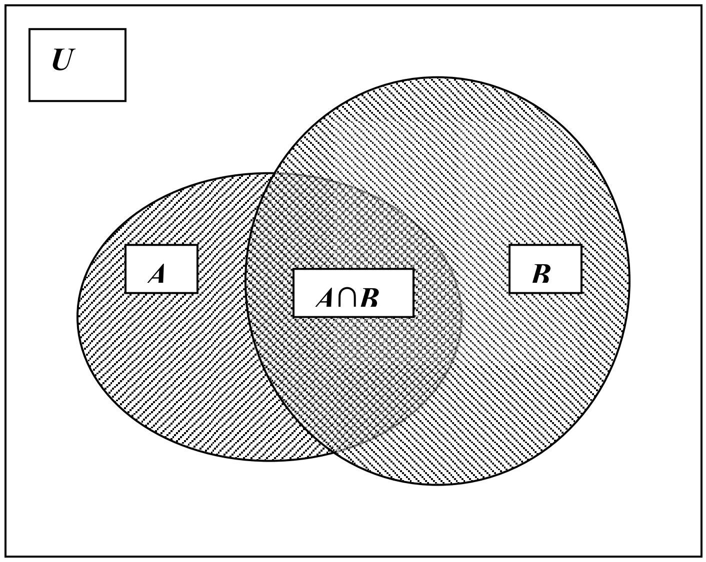

Figure 1

From this it should be clear that the probability of observing event A is true, conditional on B being true, is the ratio of the intersection to the area in B, or the first part of the equation (1)above. From symmetry, however, it should also be clear that

(2) $$\Pr\left(A \mid B\right)\Pr\left(B\right) = \Pr\left(A \cap B\right) = \Pr\left(B \mid A\right)\Pr\left(A\right)$$

Bayes Rule, as expressed in equation (1) can obtained directly from equation (2) by dividing through by Pr(*B*).

B. Classical Statistical Theory

Bayes Rule was derived within the framework of classical probability theory, which is sometimes referred to as "Frequentist" because the basic conceptual definition of probability relates to the frequency of events. For instance, what we mean by saying the probability of flipping a coin and getting heads is "one-half" (or 50%) is that in any series of trials, the ratio of the number of heads obtained to the total number of trials tends to ½ as the number of trials goes to infinity. In this context, probability can only refer to events that can be repeated (and, therefore, have some "frequency" of occurrence). Bayes Rule was derived to tell us about the conditional probability of such events, for instance the probability that it will rain conditional on it being cloudy.

Classical statistical theory is rooted in a Frequentist notion of probability. In empirical work, researchers are typically interested in some quantity (which we will refer to as θ) which they believe to have some specific (i.e., "true"), but unknown value. Examples might include the probability that a pregnant teenager will drop out of high school, or the percentage change in GDP growth that results from a certain reduction in income taxes. To form beliefs about these quantities of interest, researchers gather data and then make statistical claims by stating specific hypotheses and then assessing the likelihood of those claims being correct. For instance a researcher interested in testing whether a given coin is “fair” (that is, the probability of heads is 50%) might flip it several times and then ask whether or not the data appears to be compatible with that hypothesis.

If, however, the researcher obtains 14 heads after 20 flips, what should be concluded? With no prior information, one might believe the probability of obtaining a head was 70%. But is this what we should conclude? The classical tradition of hypothesis testing requires that a maintained (or *null*) hypothesis be clearly established and any statistical claims made by rejecting that hypothesis. Since we have good reason to assume that a coin picked at random is fair, the question we are asking is essentially: does the data reject this hypothesis? The conclusion is typically expressed as a probability: there is a 5.8% probability that we would obtain 14 or more heads in 20 flips when the probability of a head each time is 50%. This suggests that the coin may not be fair, but in making that claim there is a more than a 1 in 20 chance you will be wrong.

The key here is that the truth is considered an absolute quantity, which we cannot know but which we can make judgments about by first privileging an assumption about the truth (defining a null hypothesis); second, exploring the “sampling distribution” of the data that would result if the null were the truth; and finally, rejecting our null hypothesis only if the probability of observing the data we have seems too unlikely.

C. *Bayesian Statistical Theory*[3]

Bayesian statistics begins with the notion that Classical statistical theory has it backwards: probability in statistics is used to quantify our uncertainty, and the idea that the truth is a given and the data are probabilistic is counter to how we actually think. To see this more clearly, suppose you were asked the following question: what is the probability that Franklin Pierce was the 13th President of the United States? From a Frequentist perspective this question makes no sense at all: Pierce either was or was not the 13th President, we just may not happen to know. But at an intuitive level the question does makes sense: while those of us with an imperfect command of U.S. history may not know for sure, there is certainly a good chance that he was the 13th President (it sounds about right) and we should be able with some thought to quantify that uncertainty.[4]

Bayesian statistics is based on the idea that it is not the truth that is fixed and the data that vary, but the data that are fixed and the truth which is uncertain. In this case we can use Bayes Rule, which tells us that, given the evidence of some data *D*, our belief of the probability distribution we should assign to the range of possible values of our quantity of interest, θ, can be written

---

[3] Much of this discussion draws liberally from Leamer (1978).

[4] In fact, this happens quite often on game shows like *Who Wants to Be a Millionaire?* in which contestants effectively make monetary wagers on the outcome of their guesses about questions where they do not know the answer.

$$\Pr(\theta \mid D) = \frac{\Pr(D \mid \theta)\, p(\theta)}{\Pr(D)}$$

In the view of Bayesian statistics, such probabilities are not viewed as describing the frequency of event θ but rather as quantifying our subjective uncertainty about the (unknown) truth regarding θ. In this context, Bayes Rule gives us a clear description of how evidence can be combined with our own predispositions to arrive at own informed, yet still inherently subjective, interpretations of the "truth" regarding θ. In particular, any belief about the truth based on the data $\Pr(\theta \mid D)$ must necessarily incorporate some unconditional (or prior) belief about the data $\Pr(\theta)$. The key differences between Classical and Bayesian statistical theory are outlined in Table 1 below.

Table 1: Classical vs. Bayesian Statistical Theory

	In Classical Statistics…	In Bayesian Statistics…
the Truth	is a given	is probabilistic
the Data	is probabilistic	is a given
Probability is	Objective, representing the sampling distribution of the data under the null hypothesis	Subjective, representing our personal uncertainty about the truth
Inference from the data is obtained	by rejecting the null hypothesis	by updating prior beliefs

D. Truth, Uncertainty and Inference in a Subjective World

In the simple case where θ can take on two possible values, for instance in the logical case where a claim such as "an intelligent designer exists" is either true (T) or not true (F), we can express probabilities in terms of relative odds. Let p denote the posterior probability of the claim in question being true (that is, $p \equiv pr(T \mid D)$), and let π denote the prior probability of the claim being true. Since, the probability of the claim being false is simply one minus the probability of it being true, we can express Bayes' Rule in terms of relative odds:

(3) $$\frac{p}{1-p} = \frac{p(D \mid T)}{p(D \mid F)}\left(\frac{\pi}{1-\pi}\right)$$

Clearly, the posterior ratio of the odds is a function both of the weight of the evidence - that is the ratio of likelihoods $L(D \mid T) \equiv p(D \mid T)/p(D \mid F)$ - as well as the prior odds ratio, $O(T) = \pi/(1-\pi)$. This notation allows us to simplify the odds ratio expressed in equation (3) to simply $O(T \mid D) = L(D \mid T) * O(T)$.

III. Truth, Ideology, Bias and the Goals of Higher Education

One glaring omission in the debate over academic freedom is a clear statement of what the exact goal of higher education is, or what it should be. Most who have spent some time in academia would agree that professors are committed to the task of training the next generation, producing the most knowledgeable, productive and socially responsible future citizens that they can. The problem is that disagreements may exist over the exact criteria for each of these characteristics.[5] As a result, there has been much concern of late over what, exactly, professors are telling students in their classrooms, what their instructional goals are, and whether it is professors or students (or third-party institutions like legislatures) who have the "right" to set pedagogical parameters.

Some have, cynically, claimed that the goal of some professors is simply to "indoctrinate" students into their own particular point of view. A more generous – and I believe more realistic – interpretation of this critique, however, is that well meaning professors, aiming to convey to students their best approximation of the truth, necessarily incorporate their own prior beliefs. When those prior beliefs conflict with students' prior beliefs, students may place less weight on the information being transmitted. Kelly-Woessner and Woessner (2005) provide some empirical support for this proposition.

Since all Bayesian probabilities are interpreted as subjective beliefs, standard practice holds that the goal of one attempting convey knowledge – whether it be presenting research or teaching – should be to state the evidence embodied in $L(D|T)$ as efficiently and transparently as possible, so that readers can draw their own interpretations of the implications embodied in $\mathrm{O}(T|D)$, based on their own subjective beliefs regarding $\mathrm{O}(T)$. The practical problem, however, is that the data itself often reflects the output of researchers operating within their own statistical biases and, therefore, it is often difficult for researchers simply to aggregate, summarize or otherwise transmit to students the entire range of data in a way that is completely "value-free." By way of example, consider the question of intelligent design. It should be relative simple to survey the data available from the biology and paleontology literature and to derive likelihoods of observing such data with and without the existence of an intelligent designer. The problem, of course, is that the very question of what can or should be considered "data" is in question – many supporters and critics of intelligent design refuse to agree on what should be considered "evidence."

A more subtle version of this problem arises frequently in other disciplines. Suppose there are ten studies estimating the effect of tax cuts on growth in the U.S., five of which support the idea that tax cuts promote growth, and three of which suggest it has no effect, and two of which suggest that they slow growth. Each study employs a different methodology or dataset. How should this evidence be conveyed?

- "there is relatively little evidence that tax cuts harm growth."
- "in only one half of the studies is there evidence that tax cuts promote the growth"

[5] What, for instance, is the definition of "responsible?" Does practicing civil disobedience constitute responsible behavior? Does "most knowledgeable" prioritize scope of material or depth?

- "the evidence is mixed on the effects of tax cuts on growth, and there is no clear evidence either way."

Those ideologically predisposed to believe that taxes are intrinsically harmful to growth are likely to have an implicit bias against those studies that find that tax cuts slow growth. This may affect the way they interpret the evidence – for instance, are we really "testing a theory" or simply "measuring an impact?" A similar event occurs with the debate over the minimum wage. A generation schooled in the economic theory of price floors takes it on faith that minimum wage should have a non-negative effect on unemployment: the question is simply *by how much* unemployment will be hurt, and if it is statistically different from zero. This translates to the incorporation of "prior beliefs" regarding coefficient estimates: while the sampling distribution of the elasticity of unemployment to the minimum wage may have a familiar, symmetric "bell curve" the posterior beliefs held by researchers are asymmetric (as evidence for a negative elasticity would be treated as "no effect"). Researchers routinely dismiss coefficients that "have the wrong sign" as erroneous, and evidence of potential misspecification of their models.

Indeed the very nature of the academic process of peer review has an intrinsically social component to it. There is clearly an "optimum distance" for new research from the orthodoxy of previous research. Too little distance, and the work is considered minor or derivative; too large a distance and the work is likely to be considered flawed in some way; unbelievable. This is likely to be the greatest problem in the humanities, where the subjectivity of peer beliefs in the review process caries the greatest weight, however history has shown many examples where (supposedly objective) scientific research has run into similar problems.

B. Defining Intellectual Bias

Using the Bayesian paradigm, we have explored the nature and meaning of "truth" in a context in which the truth of any claim can never be established with certainty, and seen that our uncertainty regarding the truth is inherently subjective. We have also seen that all subjective differences can be attributed to differences in priors. Little has been said about where these priors come from, but it is reasonable to think of them as representing the outcome of a lifelong process of iterative Bayesian updating from all prior experience. In this case it is not unreasonable to claim that, from a societal standpoint, not all priors are created equally. That is, *ceteris paribus*, priors based on the cumulative lifetime experiences of a 60 year old clearly contain more information – and thus could be given more weight – than the priors formed from the lifetime experiences of a 20 year old.[6]

[6] One might, in this sense, define a "curmudgeon" as somebody with a highly concentrated prior.

C. A Formal Model of the Classroom

Define the following set of variables:

- Truth (unknown) = θ
- Evidence (large set – seen by professors) = *D*
- Evidence (small set – shown by professors to students) = *d*
- Prior of person *j* (varies – ideology, set by numerous outside factors) = $p_j(\theta)$

We can express the best estimate of the truth perceived by professors as:

(0.4)
$$E_j[\theta \mid D] = \int \theta p_j(\theta \mid D) d\theta = \int \theta \frac{p(D \mid \theta) p_j(\theta)}{p(D)} d\theta$$

In other words, the best estimate of the truth depends on the (professor's) prior beliefs about the truth. We can express the best estimate of the truth, given the full data, similarly for the student. These estimates may differ in that, given the same data, individuals with different prior beliefs may differ also in their posterior judgments.

From this various questions arise: what is the goal of professors? What should it be? What does the (vague) principle of "academic freedom" demand?

Rule 1: Professor's convey evidence in favor of what they perceive as the truth**.**

The goal of professors is typically seen, particularly by those prone to be concerned about academic "bias" as presenting students with a set of evidence (*d*) that best conveys professor's best estimate of the truth. By assumption the entire dataset (*D*) to which the professor has access cannot be conveyed completely to the student. Therefore, the professor must "select" some representative subset of data (*d*) that best conveys (approximates) the outcome perceived by the professor. Let $\mu(d)$ be the measure of the dataset *d*, and let $d \in D$.

Suppose that there is some hard constraint regarding the size of the dataset (d) that can be presented, which we will call $\bar{\mu}$. Then the professor's problem would be to choose a set of available evidence (d) of appropriate size that minimizes some loss function over the distance between the truth inferred from that data and the truth the professor believes:

(0.5)
$$\min_{\{d\}} L\left(E(\theta \mid d) - E_j(\theta \mid D)\right) \quad s.t. \quad \mu(d) \le \bar{\mu}$$

The data subset chosen by the professor in this problem explicitly reflects the students' priors. As the difference between professor and student priors grows, the data subset chosen by the professor must be increasingly skewed to "offset" student priors. This provides one pedagogical justification for a practice common among many professors of deliberately choosing to present students with evidence they know will challenge students' prior beliefs.

If students had access to the same information as

Rule 2: Professors convey only the Bayesian ratio (prior neutral)

Posterior beliefs combine prior information with evidence from the data, presented as the likelihood function $p(D \mid \theta)$. A completely "value free" presentation of the evidence to students, therefore, would dispense completely with the professor's prior information and simply present the evidence in terms of Bayes ratio.

Consider, for instance, the hypothesis that elimination of the estate tax would raise the rate of economic growth in the U.S. Let us call that hypothesis H, and its complement (that it would not raise the rate of growth) H^C.

IV. The Problem of "Perceived Bias" in the Context of Asymmetric Information

Even if students understand that the evidence *d* being presented to them is based, in part, on their professors' prior beliefs – they can use information on the distribution of those priors to make some inference regarding the large dataset (D) and the truth (given their own subjective prior).

Case 1: professors truthfully reveal their prior
Case 2: professors do not reveal their prior, and the distribution of their priors is the same as that in the population.
Case 3: professors do not reveal their prior, and the distribution of their priors differs from that of the population (the case of "institutional bias").

An often cited statistic by those on the right is that registered Democrats are over-represented in academia. The implication is that ideological *liberals* are more prevalent in the academy than in the population at large, and as a result what is being taught is likely to be "biased" in favor of liberal priors.

A. The Political Economy of "Liberal Bias"

Transparency is necessary to the proper functioning of a democracy, and our society has two traditional sources of information dissemination on public policy to provide this transparency: the media and academia. The U.S. Constitution set up an array of internal checks on power among the three branches of government; the Bill of Rights, through establishing freedom of the press, is traditionally viewed as supplementing this system with the fourth estate. the public's check on the power of the government as a whole.

Suppose, as a thought experiment, that you wanted to engage in a rent-seeking policy – defined as one that will benefit you and your interests, but at an aggregate cost to the nation at large. If this information is publicized, you may expect that majority opinion would be against adoption of your policy. One way to overcome this problem is to block or influence information flows directly; the other is to change people's perceptions of those information flows.

Institutional mechanisms, both formal such as the Bill of Rights and informal such as the tradition of Academia Freedom, make it difficult to influence information flows directly. Certainly partisan media outlets and research institutions can be created and privately funded, but to the extent that their aims are transparent and these institutions must compete with non-ideological institutions in a free market place for information, their impact is limited.

As a result, the easiest way to engage in rent-seeking policy is to change people's perceptions of the information they receive. In particular, if you want to policy legislation that promotes an agenda more "liberal" than the median voter, you would want to convince people that media outlets and academia have a conservative bias, and if you want to pass policy that promotes a conservative agenda you would want to convince people that media outlets and academia have a liberal bias. Under imperfect information, voters will examine the signals regarding policies sent by the media and academia, and will calculate their own subjective posterior probability of the policy being "bad." This probability will incorporate the information given, but filter it through their own perception of bias.

Take, for instance, the War in Iraq and let the event W represent the (true, objective) even "the war is going well" (which, for our purposes, we can define as the benefits of the War exceeding the cost), and the event W^C it's complement (that it is not going well). Although this conceptual notion may not be directly observable, a great deal of data may be observable in Iraq that will give us a sense of whether to believe W or W^C. Because the U.S. population at large cannot observe all of this data directly (that would require all 300 million people spending a great deal of time in Iraq), they rely on academics and journalists to witness this information, process it and contextualize it. In effect, the community of analysts are providing the population with a general signal about the state of the war, which we will call θ.

$$\begin{Bmatrix} W \\ W^C \end{Bmatrix} \rightarrow Data \rightarrow [academics, journalists] \rightarrow \theta \rightarrow [population]$$

If the population perceives the filtering agents as "unbiased" (meaning that their priors are the same as those of the population) then they will interpret the signal as the best (most informative) indicator of the true state of affairs. If, on the other hand, they perceive the filtering agents as biased then they must perform an induction problem.

For instance, suppose the population is given a signal that suggests that the Iraq war is a mistake (θ=M). There are four possibilities representing the cross product of two sets (the war is a mistake or not (i.e. the signal is correct), and the signal is biased or not (the signal is biased):

E1. Truth is war is not going well (signal correct), signal is not biased	(W^C, B^C)
E2. Signal incorrect, signal is not biased	(W, B^C)
E3. Signal correct), signal is biased	(W^C, B)
E4. Signal incorrect, signal is biased	(W, B)

It should be fairly obvious that whatever, the distribution of probabilities over these four possible states of the world (which must sum to 1), that the probability of the signal being "correct" (suggestive of the truth) and the probability of the signal being biased are not unrelated. Then the posterior probability one assigns to the likelihood the war is going well depends on your prior beliefs about media bias. Let b represent this prior belief regarding media bias. The Law of Total Probability can be written

$$\Pr\left(W \mid \theta = M\right) = \Pr\left(W \mid \theta = M, B\right)b + \Pr\left(W \mid \theta = M, B^C\right)\left(1-b\right)$$

Although bias may affect the signal, it is reasonable to assume that the true state of the war and the true state of media bias are uncorrelated (there is no causality between these two). This means that $\Pr\left(W \mid B\right) = P\left(W\right)$, and $\Pr\left(B \mid W\right) = \Pr\left(B\right)$. Using Bayes Rule, we can rewrite the equation above as

$$\Pr\left(W \mid \theta = M\right) = \Pr(W)\left[\frac{\Pr\left(\theta = M \mid W, B\right)}{\Pr\left(\theta = M \mid B\right)}b + \frac{\Pr\left(\theta = M \mid W, B^C\right)}{\Pr\left(\theta = M \mid B^C\right)}\left(1-b\right)\right]$$

To help make sense of this equation, suppose, for example that in the case of no bias, the probability of seeing a signal that the war is a mistake can be written

$$\Pr\left(\theta = M \mid W, B^C\right) = p$$
$$\Pr\left(\theta = M \mid W^C, B^C\right) = q$$

where p is small and q is large (so that the signals "make sense" as defined). Suppose further that the effect of (liberal) bias is to increase the probability the war is viewed as a mistake when it is not.

$$\Pr\left(\theta = M \mid W, B\right) = p + \alpha$$
$$\Pr\left(\theta = M \mid W^C, B\right) = q$$

One could certainly imagine alternative stories about bias – for instance, that it alters the signal in every state, or likelihood of a certain signal is greater in one state of the world than another – but this sort of hypothesizing is beyond the scope of this paper and since there is no reason to complicate matters unduly, we will use this as our maintained hypothesis.[7] Substituting in these probabilities, and letting π denote the prior (unconditional) probability that the war is likely to be going well, we have:

$$\Pr(W \mid \theta = M) = \pi\left[\frac{p+\alpha}{(p+\alpha)\pi + q(1-\pi)}b + \frac{p}{p\pi + q(1-\pi)}(1-b)\right]$$

If we assume "diffuse priors" about the war (that is, π = ½), then this simplifies to

$$\Pr(W \mid \theta = M) = \left[\frac{p+\alpha}{(p+q+\alpha)}b + \frac{p}{p+q}(1-b)\right]$$

If we also assume that p+q = 1, then it further simplifies to

$$\Pr(W \mid \theta = M) = \left[p + b\left(\frac{\alpha}{1+\alpha}\right)(1-p)\right]$$

The first term, *p*, represents the true posterior probability that should be assigned to the war effort going well. The remaining term is induced by the bias: it increases the posterior probability assigned to the optimistic scenario, because the media are deemed to be inappropriately pessimistic. Note, however, that as the probability *p* increases, signifying in some sense that the informativeness of the unbiased signal is decreasing, the induced bias is also reduced. Perceptions of bias is most important when the signal the population receives is understood to be most informative of the truth.

If p+q = 1, and α=0.25, then $\Pr(W \mid \theta = M) = \left[\frac{b}{5} + p\left(1 - \frac{b}{5}\right)\right] = \left[p + \frac{b}{5}(1-p)\right]$.

For example, if p= 0.1 (so q = 0.9), and α=0.25, then
$\Pr(W \mid \theta = M) = [.18b + 0.1]$

B. Beliefs about Bias and the "Salem Witch Problem"

The problem above is based on individuals having some *ex ante* belief about the likelihood of bias. The subjective probability individuals assign to such bias existing is likely to depend on the information available, however. In other words, people's belief in bias is itself subject to updating through Bayes Rule.

[7] NOTE– additive α bias affects probability calculation. Multiplicative λ bias in the case of both *p* and *q*, does NOT affect probability calculation – λ just drops out. Need to affect ratio somehow.

The main insight here is that news that suggests war not going well, against intelligent design etc. REINFORCES a belief that bias exists. If universities refuse to publish research on intelligent design, it could be because it fails to meet the standards of good science or that there is bias. If there is no bias we might or might not see such research, but if there is we certainly would not. Therefore, ceteris paribus, not seeing such research increases the posterior probability of there being bias. The only way to REDUCE such bias is to observe a university publishing research on intelligent design.

Consider, for instance, the problem of "estimating bias." Suppose again that a signal $\theta =$ M is observed, and for the sake of argument assume that it is know that the probability of such a signal being sent by the community of analysts is 0.8 if a liberal bias exists, and 0.5 if it does not. Let b once again denote the *ex ante* belief in bias existing. Figure 1 below depicts the induction problem faced by the population at large: having observed the signal "the war is a mistake" but not knowing whether or not there is bias, the population would calculate:

$$\Pr\left(B \mid \theta = M\right) = \frac{\Pr\left(\theta = M \mid B\right)\Pr\left(B\right)}{\Pr\left(\theta = M\right)} = \frac{0.8b}{0.8b + 0.5\left(1-b\right)} = \frac{0.8b}{0.5 + 0.3b}$$

As long as $b > 0$, then $\Pr\left(\theta = M \mid B\right) > b$. That is, if there is any doubt regarding the bias of the media, the presentation of evidence by the media against a hypothesis increases the subjective probability assigned to the media being biased against that hypothesis. If we denote by β the posterior probability assigned to bias existing, then the equation above suggests that:

$$\frac{\beta}{b} = \frac{8}{5+3b} \geq 1$$

Figure 1

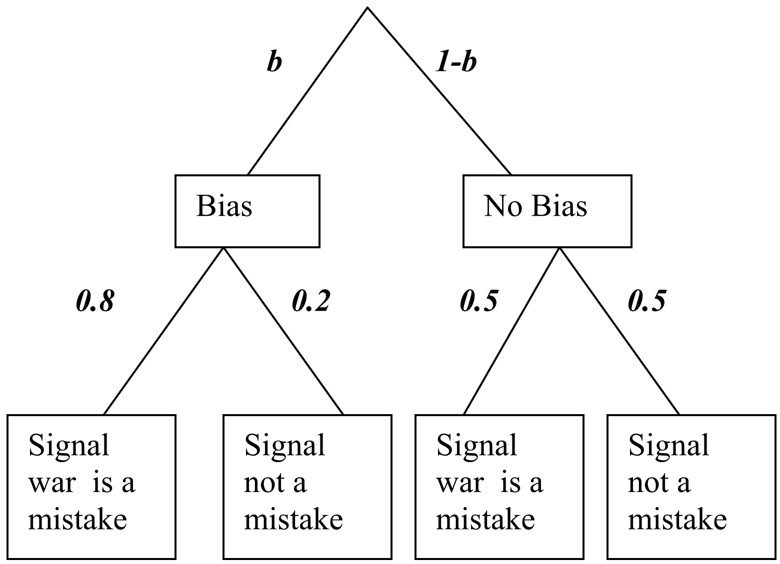

This smacks of the "Salem witch problem" – the best way to prove somebody is not a witch is to drown them and see if they survive (proof they are a witch, whereby they should be burned). The best way to prove that you are not biased against, say, the War in Iraq is to provide evidence that the War is going well – in effect, to create a reverse bias. Clearly both are inefficient when the goal is eliciting the best information possible on the true state of the world.

This "reverse biasing" effect from the asymmetric information problem is one potential reason that many of the Iraq war doves (Bush I, Powell, Kerry, Clark, etc.) had extensive military backgrounds, while the hawks (Cheney, Wolfowitz, Perle, Rice, Hilary Clinton, etc.) did not. Similarly, it is easy to imagine that those in power will always want to solidify support through generous spending initiatives regardless of ideology, but the widely accepted beliefs that Democrats favor bigger government and Republicans smaller government is likely the reason that, in recent years, both the level of government spending and the federal budget deficits (each as a share of GDP) have fallen under Democratic administrations and risen under Republican administrations.

V. Faculty Rights vs. Student Rights

Question 1: are goals opposed? Is this not just a question of achieving the greatest "surplus" in the classroom?

Question 2: can Coasian arguments be effective? What impact does assignment of liability have? Can / Does Coasian bargaining take place in the classroom? If it did, should it? Students choose professors; not vice versa.

VI. Is there a Failure in the Market for Higher Education?

One irony of the current debate over academic freedom is that the call for increased government regulation of the educational market comes from political conservatives, who are typically aligned with free-market ideologies, and opposed by political liberals who typically are more concerned about the possibility of market failures and supportive of an activist role for government in correcting them. One explanation is that the institutional structure of the academy strikes both liberals and conservatives alike as more analogous to a single monopolistic, authoritarian political institution than a system of competing firms, producing ideas and selling services in a free economic market. There is some case to be made for this view, since the institutions that would be directly affected by the legislation are those whose budgets are most heavily dependent on state subsidies. Nevertheless, these institutions are far from monopolies: even public universities are rarely monopolies in their state, and are subject to enormous competitive pressures to attract and retain talented students and faculty and to maintain enrollment numbers. In the age of the internet, information is abundant and travels quickly, so universities can scarcely afford to acquire reputations for any ideological narrowness that would compromise educational quality. Similarly, competitive pressures among institutions to produce peer-reviewed research guarantee professors "academic freedom" as well or better than any law could. What, then, is the basis for concern among the political right with respect to academic freedom: is it market efficiency, or simply equity? Put another way, is the concern of the interest groups pushing for this Bill of Rights that under the current institutional system the "pie" (the overall welfare of society) is not as large as it could be, or simply that their own share of the pie it is not as large as they would like?

The question of market efficiency is central to economic theory, and the relationship between the assumptions required to guarantee a market to delivers Pareto-efficient outcomes are well understood.[8] The notion that free markets need not, and indeed often will not, deliver equity is also well understood. The question of whether there is some failure (in the sense of optimal social welfare) of the current institutional structure of higher education thus relies heavily on two questions:

1. Can the markets for education and research be viewed like any other economic market?

2. If so, are the assumptions required for market efficiency met?

I believe the answer to both of these questions is "yes." In a competitive environment, less productive ideas, beliefs and theories – and those that hold them – should suffer the same fate as less productive firms in the U.S. economy: they will lose market share and eventually be eliminated from the intellectual marketplace. There is no guarantee of equity or "fairness" in this process, since it is one's output rather than one's input that is

[8] Efficiency is a difficult concept to define. The standard employed throughout mainstream economic analysis is Pareto efficiency (or Pareto optimality), which is a situation in which no person can be made better-off without another being made worse off.

rewarded. As a result, luck or innate talent may be rewarded as much or more than effort. These views require some examination, however.

Donald Wittman (1989) discusses the issue of whether democratic institutions are efficient. He introduces his argument by stating:

> "Many controversies in the social sciences are ultimately arguments over the nature of the market. For example, Marxist sociologists believe that both economic and political markets are characterized by poorly informed, possibly irrational, consumers and voters being exploited by monopolist suppliers of goods and policy, while (conservative) economists tend to view economic markets as working well (on the efficiency dimension) and political markets as being inefficient because of monopoly, rent seeking, and poorly informed voters."

Substituting "political" for "educational," we immediately have two potential rationales for the arguments advanced by conservative activists that educational markets may not be delivering efficient outcomes. At first blush, the language employed by Horowitz and the conservatives does suggest a Marxian paradigm –professors are "dangerous" and students must be "protected" from exploitation by the suppliers of their education. In one interview (IHE, 2006), however, Horowitz states that his purpose in writing the book *The Professors*, was "to expose the 'political corruption' of higher education," which is suggestive of rent-seeking behavior. Unfortunately, it is far from clear in the context of this interview what Horowitz means by either "political" or "corruption." The first presumably is a statement that professors' own objectives and incentives are not educational or professional, but to advance the fortunes of preferred political candidates or their polices. The second term, *corruption*, may be a reference either to uncompetitive, rent-seeking behavior or simply to a distortion of the original mission of the university. If we extract Horowitz's meaning as the presence of rent-seeking behavior by professors, then it appears his lack of faith in the educational marketplace may be likened those economists skeptical of political markets: a belief that the institutional structure of higher education institutions and/or the nature of the market itself may result in uncompetitive behavior.

This is a difficult argument to make, however. While a complete examination of the question lies outside the scope of this paper, interested readers should take note of Wittman's conclusion to the paragraph above:

> "….many of the arguments claiming that economic markets are efficient apply equally well to democratic political markets and….henceforth, the burden of proof should be on those who argue that democratic political markets are inefficient."

In other words, simply claiming that certain viewpoints are not being well represented in the academy is insufficient for restrictive government regulation over the intellectual marketplace. Some argument about market inefficiency, rooted clearly in economic theory, is necessary to justify intervention just as it would be in any other instance.

Despite the innocuousness and nearly universal appeal of "promoting intellectual diversity," the burden of proof for its necessity must be on those advancing the agenda. Economic theory offers no guarantee at all that the marketplace will deliver intellectual equity or diversity – in fact, it suggests quite the opposite since the market for ideas operates very much as a "winner-take-all" market, as modeled by Rosen (1982) and others. As a result we should expect to see certain viewpoints marginalized as a result of market forces, and appropriately so. From the standpoint of neoclassical economic theory, any attempt to enforce "intellectual diversity" in the educational market is likely to be reduce, not increase social welfare.

References

Becker, Gary (1971). *The Economics of Discrimination*, 2d ed. Chicago: University of Chicago Press.

Inside Higher Education (IHE, 2006). "David Horowitz Has a List" Internet Periodical (http://www.insidehighered.com/news/2006/02/13/list).

Students for Academic Freedom (SAF, 2006). "Academic Bill of Rights." Internet Website (http://www.studentsforacademicfreedom.org/abor.html).

Wittman, Donald (1989) "Why Democracies Produce Efficient Results," *Journal of Political Economy* 97, 1395-1424.

Kelly-Woessner, April and Matthew Woessner (2005). "My Professor is a Partisan Hack: How Perceptions of a Professor's Political Views Affect Student Course Evaluations." *PS: Political Science and Politics*. (Forthcoming)