Identity, Willful Ignorance, and the Market for News

David Bjerk
Robert Day School of Economics and Finance
Claremont McKenna College

Preliminary Draft, please do not cite

July 23, 2010

Abstract

This paper considers an environment where individuals face a claim about an issue over which they are uncertain about its truth, but must choose to act as if it is true or false each period. I show that such an environmnent can cause individuals to actively choose less-informative information sources regarding the true state of the claim, or in other words choose to be "willfully ignorant" regarding the truth. Such willful ignorance will be more prevalent when a large fraction of individuals are initially relatively unsure about the true state of the claim, and in situations where the cost to finding out one has been acting in conflict with the truth is large. The latter portion of the paper then considers how willful ignorance can affect the amount of bias in the news, revealing that increased competition in the news market has ambiguous effects on the average amount of bias in the news, but will lead to greater polarization.

In order to maintain an untenable position, you have to be actively ignorant. -Steven Colbert

1 Introduction

Around the time of the United States military invasion of Iraq in March of 2003, a CBS news poll found that 53% of the American public said they believed "Saddam Hussein was personally involved in the September 11th, 2001, terrorist attacks on the World Trade Center and the Pentagon." In the summer of 2004, after reviewing investigations by the Central Intelligence Agency, the Defense Intelligence Agency, the Federal Bureau of Investigators, and the National Security Council, as well as conducting its own investigation, the 9/11 Commission published a publicly available report (http://www.9-11commission.gov/report/911Report.pdf) that concluded that there was "no evidence indicating that Iraq cooperated with al Queda in developing or carrying out any attacks against the United States." Yet, another CBS poll taken in September 2006 found that 33% of the American public still believed that Hussein and Iraq were directly involved in the 9/11 tragedy. In April 2007, former CIA director George Tenet stated on the nationally broadcast television program *60 Minutes* that "(w)e could never verify that there was any Iraqi authority, direction, and control, complicity with al-Queda for 9/11 or any operational act against America, period." Still, a CBS poll taken in March 2008 still showed 28% of Americans continued to say that they believed Hussein was directly involved with al-Queda and 9/11.

This type of issue is certainly far more widespread than the Iraq example given above. For example, a 1997 Gallup poll showed that 44% of the American population agreed with the statement that "God created earth and man pretty much in the present form at one time within the last 10,000 years." However, there is overwhelming evidence from the scientific community that the earth is over 4.5 billion years old, and there have been hominid remains discovered that are dated as being over 4 million years old. Indeed, in the same Gallup poll, only 5% of individuals who identified themselves as scientists agreed with the statement that mankind was created by God within the last 10,000 years.

To take another example, a Florida Institute of Technology survey revealed that 24% of Americans said they believed that vaccines may cause autism and there-

fore it was safer not to have children vaccinated at all. With the exception of one medical study that appeared in the medical journal The Lancet (which has been discredited and indeed retracted by the Lancet), the primary support for such beliefs are supported by some well-read but notably non-scientific sources such as NaturalNews.com ("the internet's top natural news site"), Mothering magazine ("Inspiring natural families since 1976") and actress and former Playboy model Jenny McCarthy and her organization Generation Rescue. However, the preponderance of the scientific community, including the Centers for Disease control, the American Academy of Pediatrics, the U.S. Food and Drug Administration, and the Institute of Medicine's Immunization Safety Review Committee all have public documents stating that there is no evidence linking vaccines (or the thimerosal preservative in vaccines) to any pediatric neurologic disorders, including autism.

The first part of this paper develops a formal framework for modeling the types of phenomena highlighted above, specifically modeling why individuals may choose to stay uninformed or inform themselves with knowingly biased news and information sources regarding a claim of considerable importance to their lives, when more informative/less biased sources are readily available.

In the model, individuals are faced with a claim that is either true or false. Individuals initially do not know for sure whether the claim is true or false, but still must choose to act in a way consistent with the claim or in opposition to the claim each period, based on their beliefs regarding the veracity of the claim, as well as the benefit they will incur if it turns out they have acted in accordance with the truth, and the cost they will incur if it turns out they have acted in opposition with the truth. After choosing how to behave each period, an individual will observe further information about the true state of the claim via signals regarding its veracity. However, individuals are able to choose between different information sources, where there exists higher and lower quality information sources, where lower quality sources are less likely to reveal the truth about the claim.

The key tension in the model will be between the direct utility an individual incurs upon learning the truth about a claim and acting accordingly, and the disutility an individual would incur if he learns he has been acting in opposition to the truth, for example due to the crisis of identity he would experience due to having to renounce his past actions.

The model shows that the fraction of the population that chooses the lower-

quality information source, or in other words chooses to be *willfully ignorant,* will be greater the higher the cost individuals will incur if they find out they have been acting in opposition to the truth, as well as the more uncertain the population is about the claim. Moreover, if the cost to finding out one has been acting in opposition to the truth increases with the time one has acted in that manner, for example if an individuals increasingly identify themselves by their past behavior, whether it be psychologically or socially, then some interesting dynamics can occur. Specifically, individuals may initially choose to use a higher-quality information source regarding a new claim, but after behaving in a given way regarding that claim for a period of time, they will choose to switch to the lower-quality information source. The above results suggests that the prevalence of willful ignorance can follow a non-monotonic time trend in response to a new claim over which there is initial uncertainty over its validity. Moreover, as several simulation exercises will show, such willful ignorance can mean that even after a large fraction of the population knows the truth about a claim, a significant fraction of the population may still continue to act in opposition to that truth.

The claim that vaccines can cause autism in children is particularly illustrative of the issues to be explored in this paper. As discussed above, there is a relatively large portion of the population that believes this claim to be true in spite of the fact that the leading pediatric health organizations have widely declared that no such connection exists. So why do so many people still believe this claim is true, especially since if they knew that in actuality vaccines almost certainly do not have a connection to autism they would act differently regarding something as important as vaccinating their children? This paper argues that in situations like these, in addition to caring about and acting upon the truth, people may also want to avoid finding out that they have been acting wrong with respect to this claim because such knowledge will give them substantial disutility. For example, those who do not vaccinate their children due to concerns that such vaccines might lead to autism may come to identify themselves strongly with a "natural" living philosophy and social sphere, with the strength of this identity possibly becoming stronger the longer they continue acting in this way. If they were to learn the truth and act on it (e.g., vaccinate their children), this would cause them to have to give up that previously held identity and the psychological and social benefits enjoyed therein. Therefore, such individuals may choose to ignore or fail to consult arguably quite

reliable sources like the CDC and the American Academy of Pediatrics, and instead choose sources such as Naturalnews.com, which argues in support of the claim that vaccines can lead to autism by highlighting and putting links to other groups and individuals who claim a connection between vaccines and autism, but all of whom could arguably be termed as being part of the "natural movement." In other words, such individuals may choose to consult a news source that presents very little "new" or unbiased information regarding this debate to its readers.

The last section of the paper considers how this model of willful ignorance can inform our understanding of bias in the news market. As will be shown below, a desire by some fraction of the population to be willfully ignorant will not surprisingly lead to biased/low quality news to exist in equilibrium in the news market. More interestingly however, the model developed here shows that as the news source market becomes more competitive, for example because the costs to entering and existing in the news market fall or the relevant population of potential news consumers rises, the news market moves from an equilibrium with a somewhat biased monopoly source that everyone consumes, to a market with only high-quality/unbiased news that consumed by only a fraction of the population, to a polarized market where some individuals consume high-quality/unbiased news, while others consume news that provides very low-quality/biased news. Therefore, it is ambiguous whether, on average, consumers become more informed or less informed as a news market becomes more competitive.

2 Related Literature

This paper deals quite closely with the literature on cognitive bias, particularly the work of Rabin and Schrag (1999). They consider a model where individuals face an issue over which there is some uncertainty and update their beliefs about the "truth" based on new information signals they observe. However, key to the model is that individuals may suffer from a confirmatory bias, where it is possible that they unwittingly misinterpret a signal in opposition to their current belief/action as actually confirmatory. Rabin and Schrag show that this can lead to individuals holding incorrect beliefs even after observing an infinite amount of new information regarding the truth about the issue.

However, the psychology literature on cognitive bias suggests that not only do

people misinterpret information regarding a held belief, but just as notably, they also often actively choose less informative or knowingly biased sources of information in order to test their held beliefs. For example, individuals tend to test their beliefs about a claim by primarily considering instances that are consistent with it. One of the most cited experiments noting this phenomena is Wason's (1960) triplet-rule experiment. In it, participants were given a triplet of numbers, for example 2-4-6, and asked to hypothesize the rule that produced it. Given a triplet like 2-4-6, participants often hypothesized that the rule was "successive even numbers" and would test it by inquiring whether different sets of successive even numbers (e.g. 8-10-12, 14-16-18) also fit the rule. While such tests reveal some information, clearly they cannot rule out rules such as "increasing numbers," "numbers increasing by two," or "any three positive numbers." In other words, there are easy ways to obtain more information about the underlying rule that individuals do not choose to access. Numerous other experiments have found qualitatively similar results in other settings (Mynatt, Doherty, and Tweeney 1977, Shaklee and Fischhoff 1982, Snyder and Swann 1978, Swann, Giuliano, and Wegner 1982). Indeed, in summarizing much of this work, Nickerson (1998, p. 186) states "perhaps the confirmation bias should be thought of as a tendency to seek evidence that increases one's confidence in a hypothesis regardless of whether it should." This tendency to "seek" information that is only weakly confirmatory evidence for an individual's hypothesis is not something incorporated into Rabin and Schrag's (1999) model.

Such biased information seeking, or "willful ignorance" arises to some extent in several papers in the economics literature, however, in somewhat different contexts. For example, Benabou and Tirole (2002) develop a model where individuals choose coarser information about their ability in order to keep their self-confidence high enough to overcome their tendency to procrastinate or fail to undertake potentially beneficial actions. Like Benabou and Tirole (2002), Carrillo and Mariotti (2000) also consider a model where individuals weight current payoffs disproportionately high relative to future payoffs. Such dynamic inconsistency may mean that plans that are optimal for the current "self" may no longer be optimal for "future" selves which becomes a problem when individuals are unable to commit to a given consumption path, as the current self may foresee the future self acting in what is a suboptimal way from the current perspective. Therefore, when faced with a choice where there is some uncertainty about the outcome of that choice, individuals may choose to

forgo better information regarding the likelihood of different outcomes from that choice because the current self cannot hide that information from his future self and therefore cannot trust how that future self will act on that information. A certain amount of willful ignorance also arises in Dal Bo and Tervio's (2008) model of individual corruption. Namely, individuals may try to stay willfully ignorant of whether they are a good type or a bad type by choosing to resist temptations, even though they will immediately learn if they are bad type if they do not resist, since only bad types will succumb to temptations.[1]

Somewhat similarly to this paper, Karlsson, Loewenstein, and Seppi (2009) model what they term the "ostrich effect" in the context of investors. Specifically, under some parameterizations of their model, investors may choose to put off learning important information about their returns until a later period, even though such learning would be costless. Karlsson, Loewenstein, and Seppi's (2009) model differs substantially from the one presented in this paper however, in that their model assumes individuals incur direct utility from their beliefs about their returns at a given time (jn addition to expectations about their utility from future monetary returns), utility exhibits loss aversion and is reference dependent (where the reference point depends on the individual's initial decision regarding seeking information), and individuals receive a boost in utility when they actively seek new information.

Another closely related paper in this vein is a very interesting model by Suen (2004). In this paper, individuals are uncertain about a certain claim. Each period noisy signals regarding the veracity of the claim arise, but individuals cannot observe this new information directly, rather they must choose an intermediary (e.g. a newspaper) to "interpret" new information, where this intermediary reveals whether the emitted signal of information exceeded a given threshold and different intermediaries offer different thresholds. The model reveals that individuals will generally choose an intermediary that has a threshold biased towards the individual's prior. For example, if an individual believes a certain claim to be true, he will choose an

[1]There are also several related models where individuals do not choose to be underinformed or misinformed, but rather become misinformed due to their underlying pscyhological tendencies. For example, Blomberg and Harrington (2000) consider a model where individuals update their beliefs regarding a certain issue, but some are affected by this new information more than others. Kopczuk and Slemrod (2005) develop a model where individuals willingly repress (i.e., forget) relevant information regarding the likelihood of their own mortality in order to reduce fear of death.

intermediary that will have a threshold such that only very strong false signals will be reported. Intuitively, individuals with strong prior beliefs will only change their behavior if they observe information strongly contradicting their current action, meaning they only want a threshold such that only when the signal is drastically in opposition to their current action will it be reported.

The model developed below differs from Suen's (2004) in several ways, such as the information technology, agents being assumed to be forward looking (in contrast to Suen's model, where agents are myopic), and most notably, the differences between information sources. Specifically, in Suen, it is not so much that individuals are choosing better or worse information sources, but just different types of information. In particular, individuals face a trade off between the likelihood of finding out new information versus the quality of that information. By contrast, in the model below, there is a real quality distinction between information sources—one source has a higher likelihood of revealing the truth than the other.

Finally, this model builds on Akerlof and Kranton's (2000) conception of identity. In particular, they conceptualize a notion of identity, where individuals incur direct utility from their "identity" which is a direct function of their actions. Related to this notion of the connection between identity and action in Akerlof and Kranton's model, the model below assumes that individual's notions of identity can be tied up in their past actions, and therefore when they change actions they may lose utility due to a crises or loss of their identity.

3 Model of Willful Ignorance

Suppose each individual lives in an infinite horizon discrete time world and discounts future periods by a factor $\beta \in (0,1)$. At the beginning of time period 1, each person is faced with a claim about an issue characterized by a parameter θ that can either be "true" or "false" (e.g. "Vaccines can lead to autism" or "Under the U.S. health care reform, there will be death panels that determine end of life healthcare coverage"). However, at the beginning of period 1, an individual does not know whether the true state of θ is true or false. Rather, each person has a belief $\mu \in (0,1)$ that θ is true, where this belief is drawn randomly from a distribution F . Given his or her belief, each person must take an action $A \in \{true, false\}$, where choosing an A equal to $true$ corresponds to acting like θ is true (e.g. not vaccinating children,

lobbying against health care reform), and choosing an A equal to $false$ corresponds to acting like θ is false.

After choosing an action, an individual observes a signal regarding the true state of θ. Specifically, with probability λ an individual observes the true state of θ, and with probability $1-\lambda$ an individual observes an uninformative signal containing no new information regarding the true state of θ. If a person observes an uninformative signal, her beliefs regarding the true state of θ remain unchanged, no direct payoff is incurred, and she moves on to the next period having again to choose an action A and observing another signal regarding the true state of θ that is identical in nature to the signal observed in the previous period. In other words, when observing an uninformative signal individuals do not incur any direct utility of acting one way or the other and a person repeats the same order of events until she observes a signal that reveals the true state of θ.

Payoffs are only realized when an individual observes a signal that reveals the true state of θ. Specifically, an individual incurs a utility of $V^{T,T}$ if he learns the claim is true in actuality and the individual has acted like the claim was true. Similarly, an individual incurs a utility of $V^{F,F}$ if he learns the claim is false in actuality and the individual has acted like the claim was false. These parameters $V^{T,T}$ and $V^{F,F}$ will be referred to as the "direct utility" components of utility for a given issue.

Further, suppose that an individual incurs a disutility of $-V_t^{F,T}$ if he learns the claim is false in actuality, but the individual has acted as if it were true for the last t periods. Moreover, assume that this parameter evolves according to the function $V_t^{F,T} = V_{t-1}^{F,T} + \frac{\phi_1}{2^t}$ for some parameter $\phi_1 \geq 0$ and an initial value of $V_0^{F,T} > 0$. Note that if $\phi_1 = 0$, $V_t^{F,T}$ is simply constant over time. However, if $\phi_1 > 0$, $V_t^{F,T}$ gets larger and larger the longer an individual has acted as if θ is true, meaning the disutility a person experiences upon learning that she has been acting in in conflict with the truth about θ is greater the longer the person has been acting that way. Analogously, assume that an individual incurs a disutility of $-V_t^{T,F}$ if he learns the claim is true in actuality, but the individual has acted as if it were false for the last t periods, and again assume that this parameter evolves according to the function $V_t^{T,F} = V_{t-1}^{T,F} + \frac{\phi_2}{2^t}$ for some parameter $\phi_2 \geq 0$ and an initial value of $V_0^{T,F} > 0$.

These payoff parameters are meant to capture the key tension of the model. Namely people want to know that truth about the world as that gives them direct utility and allows them to act with certainty going forward. However, they also

have an incentive to avoid knowing the truth, as they might find out that they have been behaving in a manner in conflict with the truth, which gives them disutility. Moreover, this disutility from learning that one has been acting in conflict with the truth is allowed to increase the longer one has been acting in that manner.

This disutility from acting in conflict with the truth can arise from a variety of sources. For example, the financial loss associated with holding a flawed asset may be larger the longer one holds onto the asset as the pool of potential buyers becomes thinner and thinner. However, the possibly more interesting, and indeed preferred interpretation for the remainder of the paper is that the disutility associated with learning that one has been acting in conflict with the truth about a given claim is due to a crisis of one's identity. Specifically, with respect to some issues, an individual may come to identify him or herself by the actions the individual has taken with respect to the issue (see Akerlof and Kranton, 2000). This sense of identity with past actions may then become stronger and stronger the longer the individual has been acting in a particular way regarding this issue. This identity could be purely psychological, but it could also be social, as such a person bonds with others who believe and act similarly. Therefore, when someone learns that her past actions have actually been in conflict with the truth, she has to incur both the psychological cost of knowing she has been wrong, as well as a social cost of now being in conflict with some of her previously like-minded friends.

While not the only interpretation for the disutility incurred by learning that one has been acting in conflict with the truth regarding a given issue, for ease of explication, the disutility parameters $V_t^{T,F}$ and $V_t^{F,T}$ will be referred to as the "identity" components of utility regarding a given issue.

3.1 Choice of Action Each Period

We can first consider what action an individual should take in a given period t for an individual who believes the likelihood that $\theta = true$ is μ. Suppose he chooses to act as if the state of θ is $true$. Then with probability λ he will observe the actual state of the world, giving him a realized utility of $V^{T,T}$ if θ is actually $true$, and $V_t^{F,T}$ if θ is actually $false$. Alternatively, with probability $1 - \lambda$ he will observe an uninformative signal regarding the state of θ, meaning his belief regarding the likelihood that θ is $true$ will continue to be μ and he will be in the same place he was before he observed any signal. Therefore, given a per period discount rate of β,

it can be shown that the expected utility of an individual with beliefs μ acting as if the actual state of θ is *true* is given by

$$U_t^T(\mu) = \sum_{i=0}^{\infty}(1-\lambda)^i\beta^i\lambda\left(\mu V^{T,T} - (1-\mu)V_{t+i}^{F,T}\right).$$

Incorporating the evolution rule for $V_t^{F,T}$, the above equation can be re-written as

$$\begin{aligned} U_t^T(\mu) &= \sum_{i=0}^{\infty}(1-\lambda)^i\beta^i\lambda\mu V^{T,T} - \\ &\quad \sum_{i=0}^{\infty}(1-\lambda)^i\beta^i\lambda(1-\mu)V_t^{F,T} - \sum_{i=1}^{\infty}\sum_{k=i}^{\infty}(1-\lambda)^i\beta^i\lambda(1-\mu)\frac{\phi_1}{2^k}. \end{aligned}$$

The above expression can then be further simplified to the following closed form expression

$$U_t^T(\mu) = \frac{\lambda}{1-(1-\lambda)\beta}\left(\mu V^{T,T} - (1-\mu)(V_t^{F,T} + 2\phi_1)\right). \tag{1}$$

Using analogous reasoning, the expected utility for such an individual if he were to act as if θ is $false$ is given by

$$U_t^F(\mu) = \sum_{i=0}^{\infty}(1-\lambda)^i\beta^i\lambda\left(-\mu V_{t+i}^{T,F} + (1-\mu)V^{F,F}\right),$$

which can be solved similarly to above to obtain

$$U_t^F(\mu) = \frac{\lambda}{1-(1-\lambda)\beta}\left(-\mu(V_t^{T,F} + 2\phi_2) + (1-\mu)V^{F,F}\right). \tag{2}$$

Clearly, an expected utility maximizing individual will choose to act as if θ is *true* if and only if $U_t^T(\mu) \geq U_t^F(\mu)$, or if and only if the following expression holds.

$$\begin{aligned} &\frac{\lambda}{1-(1-\lambda)\beta}\left(\mu V^{T,T} - (1-\mu)(V_t^{F,T} + 2\phi_1)\right) - \\ &\qquad\qquad \frac{\lambda}{1-(1-\lambda)\beta}\left(-\mu(V_t^{T,F} + 2\phi_2) + (1-\mu)V^{F,F}\right) \geq 0. \end{aligned}$$

With some re-arranging, the above expression becomes

$$\frac{\mu}{1-\mu} \geq \frac{V^{F,F} + V_t^{F,T} + 2\phi_1}{V^{T,T} + V_t^{T,F} + 2\phi_2}.$$

The following proposition directly follows.

Proposition 1 *An expected utility maximizing individual will act as if θ is true at any time t if and only if his beliefs μ are such that $\mu \geq \mu_t^*$, where μ_t^* is implicitly defined by the equation below:*

$$\frac{\mu_t^*}{1-\mu_t^*} = \frac{V^{F,F} + V_t^{F,T} + 2\phi_1}{V^{T,T} + V_t^{T,F} + 2\phi_2}. \tag{3}$$

As intuition would suggest, the above proposition makes clear that those who have a higher initial belief in the likelihood that the claim is true will be more likely to act as if it is true. However, there are a couple of other things to note that come out of equation (3) in Proposition 1. First, μ_t^* is increasing in $V^{F,F}$, $V_t^{F,T}$ and ϕ_1, meaning an individual is less likely to act as if a given claim is true the higher his utility would be if he finds out the claim is false and he has acted as if it were false, and the higher his disutility would be if he finds out the claim is false but had been acting as if it were true, and the faster this disutility associated with this incorrect action increases over time. Second, μ_t^* is decreasing in $V^{T,T}, V_t^{T,F}$ and ϕ_2, meaning an individual is more likely to act as if a given claim is true the higher his utility would be if he finds out the claim is true and he has acted as if it were true, the higher his disutility would be if he acts as if the claim were false but it is actually true, and the faster this disutility associated with this incorrect action increases over time.

Finally, note that if $\mu \geq \mu_t^*$ at some period t, causing the individual to act as if the claim was true at time t, then if he does not learn the truth about the claim between period t and period $t+1$ (i.e., observes an uninformative signal), then at period $t+1$ it must be true that $V_{t+1}^{F,T} \geq V_t^{F,T}$ given the evolution function for $V_t^{F,T}$. From equation (3) in Proposition 1, it can easily be seen that this means that if $\mu \geq \mu_t^*$ at some period t, it will also be the case that $\mu \geq \mu_{t+1}^*$. In words, if an

individual acts as if the claim was true at time t and does not above the true state of the claim between period t and $t+1$, then he will also act as if the claim is true at period $t+1$. Similarly, if $\mu < \mu_t^*$ at some period t, causing the individual to act as if the claim was false at time t, then if he does not learn the truth about the claim between period t and period $t+1$, then at period $t+1$ it must be true that $V_{t+1}^{T,F} \geq V_t^{T,F}$ given the evolution function for $V_t^{T,F}$. Again, from the equation in Proposition 1, it can easily be seen that this means that if $\mu < \mu_t^*$ at some period t, it will also be the case that $\mu < \mu_{t+1}^*$. Therefore, until an individual observes the truth about a given claim, he will continue to act in the same way regarding that claim.

3.2 Choosing "Quality" of Information Regarding Claim

Now, suppose that after choosing an action regarding a given claim in a given period, each person also has to choose the "quality" of the information she will observe next period about that claim. In particular, suppose there are two information sources each individual can choose from—the "high-quality" source and the "low-quality source". The sources differ in the likelihood that they will reveal the true state of claim versus provide no new information. Specifically, the high-quality information source reveals whether the claim is true or false with probability λ_H and no new information with probability $1-\lambda_H$, while the low-quality information source reveals whether the claim is true or false with probability λ_L and no new information with probability $1-\lambda_L$, where $\lambda_H > \lambda_L$. Intuitively, the low-quality information source can either be thought of as coming from a knowingly biased news source where the bias favors the individual's initial beliefs regarding the true state of the claim, meaning it is less likely new information will be revealed, or can be thought of as an individual simply choosing to obtain less information regarding the veracity of the claim.

The key question then is who, if anybody, will choose the low-quality information source over the high-quality information source? As discussed above, this question relates to the key tension in the model, which is the desire to learn the truth and act accordingly, and wanting to continue acting the way in which one has been because one does not want to discover that one has been acting in opposition to the truth.

In order to evaluate the individual's choice of information quality let us first consider an individual who finds it optimal to act as if the claim is true, meaning

an individual whose belief regarding the claim μ exceeds μ_t^*. From above, recall that the expected utility of an optimally acting individual with such a belief is given by

$$U_t^T(\mu) = \frac{\lambda}{1-(1-\lambda)\beta}\left(\mu V^{T,T} - (1-\mu)(V_t^{F,T} + 2\phi_1)\right).$$

Taking the derivative of the above equation with respect to the quality of information parameter λ and simplifying gives

$$\frac{\partial U_t^T(\mu)}{\partial \lambda} = \frac{(1-\beta)}{(1-(1-\lambda)\beta)^2}\left(\mu V^{T,T} - (1-\mu)(V_t^{F,T} + 2\phi_1)\right).$$

It can easily be confirmed that the sign of the above expression depends on whether the sign of $\mu V^{T,T} - (1-\mu)(V_t^{F,T} + 2\phi_1)$ is positive or negative. This statement leads directly to the following key proposition of the model.

Proposition 2 *An expected utility maximizing individual who is acting as if a given claim is true will choose the low-quality information source (i.e. λ_L) at any given time t if and only if his/her beliefs at time t regarding the likelihood that the claim is true (μ) are such that $\mu < \mu_t^H$, where μ_t^H is implicitly defined by the equation below.*

$$\frac{\mu_t^H}{1-\mu_t^H} = \frac{V_t^{F,T} + 2\phi_1}{V^{T,T}}. \tag{4}$$

A very similar result can be obtained for those individuals who is choose to act as if the claim is false, meaning individuals whose belief regarding the claim μ is less than μ_t^*. From above, recall that the expected utility of an optimally acting individual with such a belief is given by

$$U_t^F(\mu) = \frac{\lambda}{1-(1-\lambda)\beta}\left(-\mu(V_t^{T,F} + 2\phi_2) + (1-\mu)V^{F,F}\right).$$

Again taking the derivative of the above equation with respect to the quality of information parameter λ and simplifying gives

$$\frac{\partial U_t^T(\mu)}{\partial \lambda} = \frac{(1-\beta)}{(1-(1-\lambda)\beta)^2}\left(-\mu(V_t^{T,F} + 2\phi_2) + (1-\mu)V^{F,F}\right).$$

Once again, the sign of the above expression depends on whether the sign of $-\mu(V_t^{T,F} + 2\phi_2) + (1-\mu)V^{F,F}$ is positive or negative, leading to the following proposition

Proposition 3 *An expected utility maximizing individual who is acting as if a given claim is false will choose the low-quality information source (i.e. λ_L) at any given time t if and only if his/her beliefs at time t regarding the likelihood that the claim is true (μ) are such that $\mu > \mu_t^L$, where μ_t^L is implicitly defined by the equation below.*

$$\frac{\mu_t^L}{1-\mu_t^L} = \frac{V^{F,F}}{V_t^{T,F} + 2\phi_2}. \tag{5}$$

Figure 1 summarizes the three propositions above. Specifically, given a fixed set of payoff parameters, individuals with beliefs regarding the likelihood a given claim is true that fall below μ_t^* will choose to act as if the claim is false, while those with beliefs that lie above μ_t^* will choose to act as if the claim is true. Moreover, the arguably more interesting result is that those whose beliefs regarding the claim that fall below μ_t^L or above μ_t^H optimally choose a high-quality information source, while those whose beliefs fall between μ_t^L and μ_t^H optimally choose a low-quality information source, on in other words, choose to be *willfully ignorant*.

Equations (4) and (5) in Propositions 2 and 3 highlight some of the key results of this model. Specifically, utility can be divided into two components: (i) the "direct" components of utility ($V^{T,T}, V^{F,F}$) which capture the value of an individual incurs from knowing he is acting in concert with the truth, and (ii) what I call the "identity" components of utility ($V_t^{F,T} + 2\phi_1, V_t^{T,F} + 2\phi_2$) which capture the feelings of regret, opportunities lost, and/or loss of identity when an individual learns he has been acting in opposition to the truth for t periods. Equations (4) and (5) show that μ_t^H will be higher, and μ_t^L will be lower, or in other words the prevalence of willful ignorance will be greater, the higher the identity components of utility are compared to the direct components.

Another important implication to come out of Propositions 2 and 3 is that once someone chooses to act in a given way, the payoffs relevant to acting in the opposite way do not affect one's choice to become willfully ignorant. For example, equation (4) shows that for a person who is acting as if a claim is true, what essentially dictates willful ignorance is relative size of the identity cost that would be incurred if he finds out that the claim is false compared to the benefit that would arise if one finds out that it was indeed true. However, the benefit that would be incurred if he had acted as if the claim was false and subsequently learned it was false has no impact on his choice over what quality of information to access. Consider this result

with respect to the vaccines and autism discussion highlighted in the introduction. While the benefit to acting as if the claim is false (i.e., vaccinating children) given it indeed proves to be false may be high (i.e., $V^{F,F}$ is large), this parameter will not affect an individual's choice to become willfully ignorant about the claim once she has made her decision not to vaccinate her children.

Finally, a further interesting implication of the model is that it suggests that people who do not start out choosing to be willfully ignorant may choose to become so over time. For example, consider an individual who is acting as if the claim is true at time t (i.e. for whom $\mu > \mu_t^*$ at time t) and has been choosing to use the high-quality information source (i.e. $\mu > \mu_t^H$). Now recall from Proposition 2 that μ_t^H was implicitly defined by the following equation:

$$\frac{\mu_t^H}{1-\mu_t^H} = \frac{V_t^{F,T} + 2\phi_1}{V^{T,T}}.$$

Note that from the assumptions about the evolution of $V_t^{F,T}$ over time, we know that if $\phi_1 > 0$, then $V_t^{F,T}$ will grow larger the longer an individual acts in the same way, which he will continue to do so as long as he does not observe a signal that informs him that the true state of θ is actually false. As can be seen in the above equation, this means that the longer the individual goes without observing the informative signal that reveals the truth regarding the claim, the higher will be μ_t^H. Therefore, after behaving the same way for a given amount of time and not observing an informative signal, μ_t^H will eventually surpass μ meaning the individual will switch to the low-quality information source. In other words, the longer an individual behaves in certain way, the more likely the individual will be to choose to become willfully ignorant. A completely analogous argument can be made for those for whom $\mu < \mu_t^L$, or those who choose to act as if the claim is false at any time t but choose (at least initially) a high-quality information source.

4 Simulating the Dynamics of the Model

This section presents the results from some simple parametric simulations of the model presented above to help illustrate basic forces at work. In all the simulations presented below, it is presumed that the claim is in actuality *true*. Moreover, the "precision" of the low-quality information source is assumed to be 0.1 (meaning with

probability 0.1 an individual will observe a signal that reveals whether or not the claim is true, and with a probability 0.9 the individual observes an uninformative signal), while the "precision" of the high-quality information source is assumed to be 0.3. Altering these parameters will not affect the basic character of the results shown below. Furthermore, the utility benefit to knowing and acting upon the truth (i.e. the *direct* component of utility) is set to be 50 , regardless of whether the claim is true or false in actuality ($V^{F,F} = V^{T,T} = 50$). The results below report the average results of ten separate simulations of populations with one hundred individuals.

4.1 Altering the Initial Distribution of Beliefs

Figures 2a and 2b show the result of two separate simulation exercises highlighting the role of the initial distribution of beliefs in determining the prevalence of willful ignorance regarding a given issue over time, as well as the fraction of the population acting in conflict with the truth regarding an issue over time. In each simulation, the *identity* component of utility is assumed to be 75 (i.e. $V_0^{F,T} = V_0^{T,F} = 75$), meaning an individual incurs a loss of 75 units of utility if he learns that his actions have been in opposition to the truth. However, in these simulations, this identity component is assumed to be fixed over time, meaning $\phi_1 = \phi_2 = 0$. Given the symmetry of the assumed parameters, individuals will choose to act as if the claim is true as long as $\mu > 0.5$ (i.e. believe that the likelihood that the claim is true is greater than one half). Moreover, given the assumed parameters, individuals will choose the low-quality information sources, or in other words choose to be willfully ignorant, if their beliefs regarding the likelihood that the claim is true are such that $0.4 < \mu < 0.6$.

Figure 2a shows fraction of individuals choosing to remain willfully ignorant (i.e., choose the lower quality information source) over time for two different initial distribution of beliefs: (i) normal with mean 0.5 and standard deviation 0.15 (solid line), and (ii) normal with mean 0.5 and standard deviation 0.05 (dotted line). As can be seen, the fraction of the population choosing to be willfully ignorant each period is higher under the distribution with less variance (i.e. standard deviation of 0.05). In words, Figure 1a illustrates that there is greater willful ignorance when there is less certainty regarding whether initial actions correspond to the truth. Indeed, when the initial beliefs are drawn from the higher variance distribution (where there are more people relatively certain that their actions correspond to the truth), only about

20% of the population are using the poor-quality information source by period 10. By contrast, when the initial beliefs are drawn from a distribution with relatively little variation centered around 0.5 (meaning most people are very uncertain about whether their initial actions actually correspond to the truth), almost 40% of the population is still choosing poor-quality information by period 10.

These differences in the extent of willful ignorance across simulations manifest themselves in the evolution of actions over time. Specifically, Figure 2b shows how the average fraction of individuals acting in opposition to the truth changes over time under the two different distributions of initial beliefs. As can be seen, there is a higher fraction of individuals who continue to act in opposition to the truth each period under the distribution of initial beliefs with the lower variance.

It is worth noting that the downward time trends in willful ignorance and acting in opposition to the truth are due to the fact that even when individuals choose to stay willfully ignorant (i.e., use the low-quality information source), there is still some chance of learning the truth about the claim each period (in this case a likelihood of 0.10). As discussed above, once they learn the truth they act in accordance with it and they are assumed to choose the high-quality information source from there on out. If the low-quality information source was assumed to be even worse (e.g. reveal the truth with a likelihood of 0.05 or 0.01 each period), the fraction of individuals acting in opposition to the truth each period would be even larger.

In summary, the simulations in Figures 2a and 2b reveal that according the model, the prevalence of willful ignorance, and the corresponding prevalence of actions in opposition to the truth, will be higher the more uncertain the population is initially regarding the truth about the claim. Intuitively, when people are relatively uncertain regarding whether or not their actions correspond to the truth, and if they pay a relatively large cost if they learn that their actions have been in opposition to the truth, they may be more likely to try to avoid such a shock by remaining willfully ignorant even if that means acting in opposition to the truth.

4.2 Altering the Identity Component of Utility

As alluded to above, one of the more interesting aspects of this model is to consider the implications when the identity component of utility grows over time, meaning when $\phi_1, \phi_2 > 0$. In particular, Figure 3a shows the evolution of willful ignorance in

a population where initially individuals only care about acting in accordance with the truth, but over time start to incur a greater and greater cost if they find out that they have been acting in opposition to the truth. In particular, this simulation assumes $V_0^{F,T} = V_0^{T,F} \approx 0$, but $\phi_1 = \phi_2 = 25$. This simulation assumes initial beliefs are distributed normally with a mean of 0.5 and a standard deviation of 0.10 (an amalgam of the two belief distributions used in Figures 2a and 2b).

As can be seen in Figure 3a, given these assumptions, willful ignorance evolves according to a non-monotonic pattern. In the first couple of periods very few individuals choose to be willfully ignorant since they care only about learning about and acting in accordance with the truth. However, after behaving in a given way for a couple of periods, they start to identify with their past behavior, or otherwise start to feel that they will incur a good deal of disutility if they learn that they past actions have been incorrect. This causes individuals to start switching to the low-quality information source over time, causing willful ignorance to increase. After awhile however, more and more individuals (even those choosing to be willfully ignorant) start to learn the truth, causing willful ignorance to start to fall again (again assuming willful ignorance is not possible if the truth is known).

Figure 3b shows how such willful ignorance affects behavior over time. Specifically, the solid line in Figure 3b shows how the fraction of individuals acting in opposition to the truth falls over time when only the high-quality information source is available (i.e., a world without the possibility of willful ignorance). The dashed line shows how shows how the fraction of individuals acting in opposition to the truth evolves when individuals are allowed to switch over to the low-quality source. Not surprisingly, allowing individuals to be willfully ignorant drastically increases the fraction of individuals acting in opposition to the truth several periods out.

Clearly, altering the precision of the information sources will alter the dynamics in all of the simulations discussed above, but the general patterns will stay the same in each. These simulations illustrate that when most of the population is initially very uncertain about the truth of a given claim, a relatively large fraction of the population may choose to remain willfully ignorant regarding the truth of that claim and continue to act in opposition to the truth for substantial periods of time. This will be true even if individuals start out caring only about learning and acting upon the truth.

The one parameter change that would have a qualitative impact on results would

be if the low-quality signal was assumed to be pure noise, meaning it had no chance of revealing the truth about the claim (i.e. $\lambda_L = 0$). This can be thought of as ignoring any new information. In that case, individuals would still choose such an information source if their prior was sufficiently far from certainty (i.e., sufficiently far from zero or one) and the identity component of utility was strong enough. However, in contrast to the above simulations, such individuals would then never learn the truth, never switch from the low-quality information source (in this case pure noise), and therefore always continue to act in opposition to the truth. Such a situation is most interesting with respect to the simulation where initially individuals only care about the truth, but come to identify with their past actions the longer they have acted that way. In this case, even though everyone initially cares only about learning and acting in accordance with the truth, over time many individuals will stop paying attention to real information, and thereby hold on to their incorrect beliefs and act in opposition to the truth forever.

5 Discussion of Basic Model

While relatively simple, this basic model can provide a good deal of perspective on many different seemingly incongruous situations. Considering again the examples from the introduction, in all of those cases, it is clear that for some people there is a substantial amount of what I would term identity caught up in the issue at hand. Those who say they still believe Saddam Hussein was behind 9/11 likely strongly advocated for the initial invasion and strongly identify with the neoconservative movement. Similarly, there is no question that the vast majority of individuals who say they believe God created earth and man within the last 10,000 years strongly identify themselves as religious Christians, an identity which influences a good part of their actions and social lives.

Returning to the vaccines and autism claim, it is informative to consider the available information regarding this claim. The most common claim in this literature deals with the MMR (Measles, Mumps, and Rubella) Vaccine. Dozens, if not hundreds of peer-reviewed studies, in the top medical journals such as the *Journal of the American Medical Association* and the *New England Journal of Medicine* have found no evidence of a causal or even correlative relationship between childhood vaccines and autism. Indeed, the Institute of Medicine of the National Academy

of Sciences has stated "(t)he body of epidemiological evidence favors rejection of a causal relationship between the MMR vaccine and autism spectrum disorder" (Stratton et al., 2001), and lists the voluminous peer reviewed scientific studies on which it bases this conclusion. Moreover, this information is very accessible, for example under "Common Questions Parents Ask about Infant Vaccinations" on the Center For Disease Control's website, it explicitly states "Scientific studies and reviews have found no relationship between vaccines and autism."

As stated in the introduction, despite the above conclusions from arguably quite high quality information sources, a substantial number of individuals state that they believe that vaccines cause autism and follow through on such stated beliefs by not vaccinating their children. One of the most visible advocates for the claim that vaccines can cause autism has been Generation Rescue. The "Science" section of the the Generation Rescue website offers citations for a variety of claims regarding vaccines and neurological disorders such as autism. What is interesting to note is that most of the citations listed are from high-quality scientific peer reviewed journals, the same as those used by the Institute of Medicine, suggesting that Generation Rescue considers these sources to be of high-quality. However, regarding the key claim that "Vaccinated children have higher rates of autism and ADHD than unvaccinated children" the only listed source for this claim is an unpublished phone survey conducted by Generation Rescue. It seems hard to argue that such a survey could be of similar quality to the many peer reviewed scientific studies cited by the Institute of Medicine and the CDC. Indeed, in an interview on the television show Frontline, J.B. Handley, the founder of Generation Rescue, arguably suggests his unwillingness to evaluate the broader scientific evidence regarding this issue when, referring to the onset of autism in his son, he states "(y)ou don't have any science that can show me that the regression wasn't triggered by the six vaccines." (Frontline, 2010)

There are a variety of other situations where this model of willful ignorance can provide some perspective. Consider simple Principle/Agent problems. Suppose an agent sees his reputation based on what he has promoted regarding a given claim. For example, consider an investment advisor who advocated his clients to invest in mortgage backed securities. Given he will lose some of his reputation if he learns that he gave bad advice, such an agent may choose to be willfully ignorant and ignore information that may reveal that his advice to his client has been incorrect.

Similarly, consider group based discrimination in hiring. Employers may choose

not to hire individuals of a certain race or gender group, not because they have an active distaste for such individuals, but because they truly believe that there is a high likelihood that such individuals do not possess the necessary skills to do the job correctly. Because even such employers may want to identify themselves as not being racist or sexist, they may persist in this discriminatory behavior, and even avoid new information regarding the actual skill of such individuals, in order to avoid learning that their behavior was racist or sexist in the sense that such individuals could successfully do the job but were be excluded solely due to their skin color or gender.

Another historical example relates to the relationship between members of the American Communist Party and the Soviet Union. In particular, in 1930s United States, a substantial number of left-leaning American intellectuals who identified with a socialist philosophy believed in, actively supported, and often acted at the behest of the Joseph Stalin and the Soviet Communist party, engaging in activities ranging from organized letter writing campaigns to spying. This behavior all occurred despite the political purges, economic shortages, Gulags, and overall curtailment of freedom that was occurring in the Soviet Union at the time. In his autobiography, philosopher Sidney Hook states that one reason why he and some of his colleagues committed such errors of judgement was "genuine ignorance of what was really occurring within the Soviet Union. Perhaps this was exacerbated by a reluctance to find out." He goes on to say, "(s)ome who gradually came to agree (that the Soviet Communist party had betrayed many of the socialist ideals) confessed that it required more courage to break with the communist movement than to join it. They had made so much of an emotional investment in it." (Hook, 1987 p. 175-177).

One thing these examples highlight, is that the consequeces of willful ignorance often fall on others.

6 Willful Ignorance and the Market for News

There has been much written in both the popular press, and more recently in the economics literature, about biased news sources. For example, Baron (2006) considers a supply side explanation, where personal preferences or career concerns can cause journalists to choose to provide biased coverage of an issue. More closely related to

the model developed here are the demand side explanations posed by Mullainathan and Shleifer (2003) and Gentzkow and Shapiro (2006). In Mullainathan and Shleifer (2005), there exist individuals who are assumed to want to know the truth and individuals who are assumed to prefer to hear news consistent with their own prior beliefs. The presence of these news consumers with a preference for bias can cause news sources to offer only biased news in the case of a monopoly, or polarizing viewpoints in a duopoly setting. In Gentzkow and Shapiro (2006), news sources want to build a reputation for quality. However, rational consumers will tend to place more faith in news sources whose reports are generally consistent with their prior beliefs. This causes a demand for biased news to arise endogenously, causing news sources to slant their coverage to some degree, especially concerning issues that will take a long time to reveal their true state.

As will be shown below, the forces in the model developed in this paper have quite straightforward implications for the demand for biased/low-quality news. While some of the results are similar to Gentzkow and Shapiro (2006) and especially Mullainathan and Shleifer (2003), they have a slightly different flavor and actually combine, and in some sense extend, the insights of these papers.

6.1 News Source Covering Multiple Issues

To keep the model as simple as possible, let us assume that there is a fixed cost of c for a firm to enter and stay in the news source business each period, and revenue for a news source is simply some parameter R times the number of individuals using that source each period. For example, R can be thought of as advertising revenue that depends on number of consumers. Furthermore, suppose there are $m > 1$ different claims regarding various issues of interest. In such a world, let us denote a given news source k simply as a vector $S^k = [\lambda_1^k, \lambda_2^k, ..., \lambda_m^k]$, where each λ_i^k is the quality of information that news source k provides regarding issue i (i.e., λ_i^k is the likelihood of revealing the truth about issue i any given period).

Furthermore, assume the highest possible quality of information a news source is able to provide regarding issue i is $\overline{\lambda}_i \in (0, 1]$. Alternatively, suppose that by keeping completely uninformed, each individual has a likelihood of $\underline{\lambda}_i \in [0, \overline{\lambda}_i)$ of learning the truth about issue i any given period. This means that for each claim i, a news source can pick any λ_i between $\underline{\lambda}_i$ and $\overline{\lambda}_i$. A news source that does not choose to set a given λ_i equal to $\overline{\lambda}_i$ (i.e., to provide the highest possible quality news

regarding claim i) can be interpreted in two ways. On the one hand, we can think of that news source as simply providing relatively more limited information regarding that claim than is possible. On the other hand, we can interpret such a news source as simply providing meaningless information that either supports the claim as being true or false. Either way, given choice of λ_i does not affect a news source's costs, we can think of a news source that chooses a $\lambda_i < \overline{\lambda}_i$ as providing *biased* coverage of issue i, in the sense that they could present more or better information regarding claim i but choose not to.

6.2 Individual Utility Under a Given News Source

Now consider an individual for whom $\mu_i > \mu_i^*$ for claims 1 through m/κ (for $\kappa \in [1, m]$), but $\mu_i < \mu_i^*$ for claims m/κ through m. In other words, consider an individual who finds it optimal to act as if claims 1 through m/κ are true, but finds it optimal to act as if claims m/κ through m are false. If we then label an individual's beliefs and preferences regarding a given claim i with a subscript i, then from equations (1) and (2), the individual's expected overall utility at any given time t from consuming a news source $S^k = [\lambda_1^k, \lambda_2^k, ..., \lambda_m^k]$ is given by

$$\begin{aligned} W_t(\lambda_1^k, \lambda_2^k, ..., \lambda_m^k) &= \sum_{i=1}^{m/\kappa} \frac{\lambda_i^k}{1-(1-\lambda_i^k)\beta}\left(\mu_i V_i^{T,T} - (1-\mu_i)(V_{i,t}^{F,T} + 2\phi_{i,1})\right) + \\ &\quad \sum_{i=m/\kappa}^{m} \frac{\lambda_i^k}{1-(1-\lambda_i^k)\beta}\left(-\mu_i(V_{i,t}^{T,F} + 2\phi_{i,2}) + (1-\mu_i)V_i^{F,F}\right). \end{aligned}$$

From the above equation and Propositions 2 and 3, we know that an individual's expected overall utility will be maximized by a news source $S_k = [\lambda_1, \lambda_2, ..., \lambda_m]$ such that if $\mu_i V_i^{T,T} - (1-\mu_i)(V_{i,t}^{F,T} + 2\phi_{i,1}) < 0$ or $-\mu(V_t^{T,F} + 2\phi_2) + (1-\mu)V^{F,F} < 0$ (or equivalently if $\mu_{i,t}^* < \mu_i < \mu_{i,t}^h$ or $\mu_{i,t}^\ell < \mu_i < \mu_i^*$) then $\lambda_i^k = 0$, and if $\mu_i V_i^{T,T} - (1-\mu_i)(V_{i,t}^{F,T} + 2\phi_{i,1}) > 0$ or $-\mu(V_t^{T,F} + 2\phi_2) + (1-\mu)V^{F,F} > 0$ (or equivalently if $\mu_i > \mu_{i,t}^h$ or $\mu_i < \mu_{i,t}^\ell$) then $\lambda_i^k = 1$. In words, for each claim, individuals would *optimally* choose a news source that gives them no information about any claims over which they prefer to remain willfully ignorant, but fully informs them about any claims over they do not wish to be willfully ignorant.

Finally, assume individuals will consume the news source that gives them the

highest expected utility possible, subject to the requirement that the news source gives them at least the same expected utility as consuming no news source, or equivalently the same expected utility as a news source defined by the quality vector $[\underline{\lambda}_1, \underline{\lambda}_2, ..., \underline{\lambda}_m]$.

6.3 News Sources in Equilibrium

The basic question to be addressed here is what types of news sources can exist in an equilibrium? In general, the answer to this question will depend on the distribution of preferences and beliefs in the population. However, the simple environment developed below will be used to illustrate some interesting outcomes that can arise in equilibrium.

Again to keep things as simple as possible, suppose that there are two issues/claims of interest, where issue 1 is an issue with a very small or negligible identity component, while issue 2 has a stronger identity component. For example, suppose issue 1 corresponds to a claim such as "this winter will be unseasonably cold" or "Julia Roberts likes Starbucks lattes," while issue 2 corresponds to a claim such as "the healthcare reform bill will require a 20 percent tax increase of the next decade" or "if CO_2 emissions stay at the current level or higher there will be large scale floods of coastal communities in the next 20 years."

Next assume there are two types of individuals. First, there are N_1 type-1 individuals, who are such that they want the best information possible on all issues, meaning $W_t(\overline{\lambda}_1, \overline{\lambda}_2|type-1) \geq W_t(\lambda_1, \lambda_2|type-1)$ for all $\lambda_1 \in [\underline{\lambda}_1, \overline{\lambda}_1]$ and $\lambda_2 \in [\underline{\lambda}_2, \overline{\lambda}_2]$. Second, there are N_2 type-2 individuals, who are such that they want the best information possible on issue 1, but want to remain willfully ignorant on issue 2, meaning $W_t(\overline{\lambda}_1, \underline{\lambda}_2|type-2) \geq W_t(\lambda_1, \lambda_2|type-2)$ for all $\lambda_1 \in [\underline{\lambda}_1, \overline{\lambda}_1]$ and $\lambda_2 \in [\underline{\lambda}_2, \overline{\lambda}_2]$.

In order to see how simply the presence of willfully ignorant news consumers can affect equilibrium outcomes, we can focus on the case where $N_1 > N_2$, or where a majority of individuals still want the highest quality news possible on all issues. Finally, let us further limit our analysis to only individuals whose beliefs regarding issue 2 are such that $\mu_2 > \mu^*_{2,t}$, meaning they all choose to act as if the claim regarding issue 2 is true. As will be seen below, all of the following analysis will identically apply to those who find it optimal to act as if claim 2 is false.

Let us first consider a situation where the costs to entering and existing in the

news source business are large enough and/or where the potential market of consumers is small enough such that $(N_1+N_2)R > c > N_1R > N_2R$. In words, consider a situation where in order to cover its costs, a news source must capture members of both types. In order to proceed, define $\widetilde{\lambda}_2$ to be such that $W_t(\overline{\lambda}_1, \widetilde{\lambda}_2|type-2) = W_t(\underline{\lambda}_1, \underline{\lambda}_2|type-2)$, meaning $\widetilde{\lambda}_2$ is the quality parameter for claim 2 such that type-2 individuals are indifferent between consuming a news source with a quality vector $[\overline{\lambda}_1, \widetilde{\lambda}_2]$ and consuming the lowest possible quality news vector $[\underline{\lambda}_1, \underline{\lambda}_2]$ (or equivalently consuming no news). We will assume $\widetilde{\lambda}_2 < \overline{\lambda}_2$, meaning $type-2$ individuals prefer no news to a vector giving the highest quality information possible. Given this definition, we can now state the following proposition.

Proposition 4 *If $(N_1+N_2)R > c > N_1R > N_2R$, then all Nash Equilibria are such that there exists one news source offering a quality vector $[\overline{\lambda}_1, \lambda_2^*]$, where $\lambda_2^* \in [\underline{\lambda}_2, \widetilde{\lambda}_2]$, and all types consume this news source.*

The intuition for this proposition is quite straightforward. When there are high costs to entering and existing in the news source market, or the potential market of news consumers is relatively small, only one news source can exist, where this news source provides good information on non-identity type issues, but little or relatively biased/uninformative information about identity type issues in order to ensure that those who wish to remain willfully ignorant on those issues will still choose to consume news. As long as the firm in the market offers a news quality for issue 2 that is equal to or lower than $\widetilde{\lambda}_2$, there is no incentive for another news source to enter, as they will not be able to attract enough consumers regardless of what news quality vector they offer. It is interesting to point out however, that a whole continuum of different news qualities with respect to issue 2 can be offered by the one news source in equilibrium, from completely uninformative news about issue 2 (i.e., $\underline{\lambda}_2$), to biased but relatively informative news about issue 2 ($\widetilde{\lambda}_2$). The firm is simply constrained from offering such high enough quality news regarding issue 2 that those who would prefer to be willfully ignorant regarding issue 2 would rather consume no news than the news being offered.

Importantly, note that this proposition implies that under these circumstances, even those individuals who do not want to remain willfully ignorant on any issue will be constrained by the marketplace to consume a less than fully informative news source.

Next, let us consider the case where the cost of entering and existing in the news business is somewhat smaller and/or the population of non-willfully ignorant potential consumers is large enough such that $N_1R > c > N_2R$, meaning a news source can exist even if it only sells to type-1 individuals. This lead to the following Proposition.

Proposition 5 *If $N_1R > c > N_2R$, then in any Nash Equilibrium, news sources will always offer the highest possible quality vector $[\overline{\lambda}_1, \overline{\lambda}_2]$, but only type-1 individuals will consume any news.*

The intuition for this proposition is slightly more complicated. First note that it cannot be the case that in equilibrium there is news source that offers a news quality vector less than $[\overline{\lambda}_1, \overline{\lambda}_2]$. To see why, suppose it was the only news source. In this case, it would always make sense for an entrant to enter and offer $[\overline{\lambda}_1, \overline{\lambda}_2]$, since by doing so the entering firm could get all of the type-1s, which would assure a positive profit, while the incumbent would get only type-2s, of which there is not enough of to make a profit given the costs of staying in business. Moreover, as long as there is at least one high-quality news source (i.e., offering the quality vector $[\overline{\lambda}_1, \overline{\lambda}_2]$), any firm offering a lower quality vector than $[\overline{\lambda}_1, \overline{\lambda}_2]$ will get only type-2s, of which there are not enough of to ensure a positive profit given the costs of entry. Therefore, there can only exist news sources offering $[\overline{\lambda}_1, \overline{\lambda}_2]$ in equilibrium. If we then assume that if there are multiple news sources offering the identical quality vectors they split the market equally, then in equilibrium there will exist s_1 such high-quality news sources, where s_1 is the largest integer such that $N_1R/s_1 \geq c$.

Finally, suppose the costs to entering and existing in the news business are small enough and/or population of potential consumers is large enough such that $N_1R > N_2R > c$, meaning a news source can operate even if it is only capturing members of one type. Under these circumstances, the following Proposition holds.

Proposition 6 *If $N_1R > N_2R > c$, then in any Nash Equilibrium there will exist two kinds of news sources—general high-quality sources who offer the highest possible quality vector $[\overline{\lambda}_1, \overline{\lambda}_2]$, and specific/biased sources who offer the quality vector $[\overline{\lambda}_1, \underline{\lambda}_2]$, with all type-1 individuals consuming a high-quality news source and all type-2 individuals consuming a specific/biased source.*

Again, the intuition for this proposition is quite straightforward. The market is big enough/costs of entry low enough such that news sources can specialize by targeting one type of consumer. Note that if a news source tried to offer something other than $[\overline{\lambda}_1, \overline{\lambda}_2]$ or $[\overline{\lambda}_1, \underline{\lambda}_2]$ it would not attract any customers. If we again assume that if there are multiple news sources offering the identical quality vectors they split the market equally, then in equilibrium there will exist s_1 such high-quality/unbiased news sources, where s_1 is the largest integer such that $N_1 R/s_1 \geq c$, and there will exist s_2 such low-quality/biased news sources, where s_2 is the largest integer such that $N_2 R/s_2 \geq c$.

In words, Proposition 6 says that if the costs of entering and existing in the news business are low enough, or if the relevant population is large enough, there will exist both uniformly high-quality news sources, who attempt to provide unbiased high-quality news on a broad array of issues, as well as specific/biased news sources who offer high-quality information on non-identity type issues, but offer little or biased/uninformative coverage of other issues that contain a strong identity component.

While this two issue, two type model is admittedly quite simple, it still highlights some interesting issues. Comparing the equilibrium outcomes across propositions, we can see that as the cost of entering and existing the news business falls, or the size of the relevant consumer population increases, the number of news sources expands. Moreover, while the average quality of news/amount of bias can go up, down, or stay the same (depending on whether λ_2^* is greater than, equal to, or less than $\frac{N_1\overline{\lambda}_2+N_2\underline{\lambda}_2}{N_1+N_2}$), there will exist market polarization, with both high-quality/unbiased news sources and low-quality/biased news sources.

Perhaps even more interesting is to consider what happens to the amount of ignorance in the population as c falls (or the size of the relevant population increases). Those individuals who do not want to remain willfully ignorant on any issues (type-1 individuals) will be able to find a high-quality news source, which decreases the amount of ignorance among this group. However, this comes at a cost. Specifically, as c falls, type-2 individuals will either consume no news, or at best consume much more biased news than they were initially, making them more ignorant than they would be in a world with higher entry costs and less competition in the news market.

It is informative to compare the results coming from the propositions above to those arising in some the models of media bias discussed previously. For example,

like Gentzkow and Shapiro (2006), this model predicts less bias in the news coverage of some issues than others. In this case, those issues where individuals do not develop strong identities based on their previous actions, such as analysis of sporting events, entertainment, weather and stock market trends, will always enjoy relatively high quality/unbiased coverage. On the other hand, those issues where some individuals are prone to develop strong identities based on their previous actions, for example debates about the success of an armed conflict, regulation, and religion, may incur relatively low quality/biased coverage. While Gentzkow and Shapiro obtain a similar result, it is for very different reasons. Specifically, issues that reveal their true state quite quickly, particular sporting events, today's weather and stock market activity, will see little bias, since those who provide biased coverage will quickly be revealed to be providing lower quality/biased news. Alternatively, issues where the true state is revealed only very gradually, the success of an armed conflict, the costs of a given regulation, the truth about morality in the eyes of God, will see substantial bias since those who provide biased coverage may be embraced by a market place who wants to believe they are correct and place more confidence in news sources that confirm their prior beliefs.

The above results also relate to a discussion running through both Mullainathan and Shleifer (2005) and Gentzkow and Shapiro (2006) regarding television news in the Arab world. Gentzkow and Shapiro's model implicitly suggests that subsidizing entry of new ostensibly unbiased news sources to compete with al-Jazeera can lead to a more accurately informed Arab population, as the competition posed by such a new news source could have a moderating effect on bias in al-Jazeera's coverage of the western world and its relationship to the Arab world. Alternatively, Mullainathan and Shleifer's model suggests that entry by a less-biased news source may cause an entity like al-Jazeera to move toward even more bias to differentiate itself in order to maintain market share.

One could argue that the implications of the model developed here captures both of these forces. Specifically, as costs of entry fall and the news market becomes more competitive in a given locale, the market moves from one source offering moderately biased coverage, to a more polarized news market, with some news sources offering high quality/less biased news, while others offer low quality/very biased news (at least on issues with a strong identity component). So lowering entry barriers can simultaneously lead to both a proliferation of more informative news sources, as well

as a greater radicalization of news sources. The effect of this is as stated above—some individuals will now be able to act on their desire to better inform themselves, while others will be better able to maintain their willful ignorance on the more controversial identity fraught issues. The net affect of easing market entry conditions on the overall quality of information consumed by the population is ambiguous. However, given those who wish to remain willfully ignorant are arguably the ones some policy makers most want to better inform, this model suggests it might not be as simple as facilitating more competition in the news business.

7 Summary and Conclusion

Why would someone want to remain willfully ignorant regarding a particular claim? As this paper argues, individuals may prefer willful ignorance to better information regarding the claim when they are initially relatively uncertain about the truth of the claim and when they foresee that they will incur substantial disutility if they learn they have been acting in conflict with the truth about the claim. The question then is what types of issues may cause individuals to feel they will incur such disutility from learning they have been acting in conflict with the truth? As discussed in the text above, substantial disutility from learning one has been acting in conflict with the truth is likely to arise in issues where people start to identify themselves strongly with their past behavior.

As the relative simplicity of model developed above shows, willfully ignorant behavior can arise even under modest assumptions regarding payoffs. Not surprisingly, given the ability to be willfully ignorant, relatively large fractions of the population can have beliefs and act in ways that are almost incomprehensible to the rest of the population. More notably, when a claim regarding a new issue arises (e.g., hormone therapy is a relatively safe way for women to treat physiological changes during menopause), willful ignorance may follow a non-monotonic pattern. Specifically, individuals may initially care only about learning the truth and acting correctly, and therefore try to obtain the most informative information possible regarding this issue. However, over time, as they have been acting one way or the other about the claim but still uncertain regarding the truth about the claim, individuals may start to develop a strong identity attachment to their previous behavior. This may cause them to stop seeking new information, or simply look for information supporting

their past behavior.

The model also showed how such willful ignorance can influence the type of news being offered in the marketplace. Most notably, the model reveals that there is likely somewhat of a trade-off in terms of the quality of the news being reported as a news market becomes more competitive. Specifically, in small markets or markets where there are high costs to entry, the news market may consist of a single source offering somewhat biased news. While this biased news monopoly may impede those who wish to be more fully informed, it also can mean that those who wish to remain willfully ignorant on some issues are unable to avoid at least consuming some information on those issues. By contrast, as the market expands or the costs to entry fall, high quality news operations will be able to exist and provide unbiased news to those looking to stay well informed. However, this comes at the cost of those wishing to remain willfully ignorant on some issues to either consume very biased news on those issues, or consume no new information on all issues. In the end, this suggests that if one's goal is to ensure a better informed public, the key may not lie in expanding the media market, but rather in convincing people that it is okay to change one's mind and behavior. In the words of the 2004 United States presidential election, being a "flip-flopper" should be seen as something to be encouraged rather than mocked (at least in some contexts).

References

[1] Akerlof, George and Rachel Kranton. (2000). "Economics and Identity." Quarterly Journal of Economics 115(3): 715-733.

[2] Baron, David. (2006). "Persistent Media Bias. "Journal of Public Economics 90: 1-36.

[3] Benabou, Roland and Jean Tirole. (2002). "Self-Confidence and Personal Motivation." Quarterly Journal of Economics 117(3): 871-915.

[4] Blomberg, S. Brock, and Joseph Harrington. (2000). "A Theory of Flexible Moderates and Rigid Extremists with an Application to U.S. Congress." American Economic Review 90: 605-620.

[5] Carrillo, Juan D. and Thomas Mariotti. (2000). "Strategic Ignorance as a Self-Disciplining Device." Review of Economic Studies 67: 529-544.

[6] Dal Bo, Ernesto and Marco Tervio. (2008). "Self-Esteem, Moral Capital, and Wrongdoing." NBER Working Paper 14508.

[7] Frontline. (2010). "The Vaccine War." April 27, 2010.

[8] Gentzkow, Matthew and Jesse M. Shapiro. (2006). "Media Bias and Reputation." Journal of Political Economy 114(2): 280-316.

[9] Hook, Sidney. (1987). Out of Step: An Unquiet Life in the 20th Century. Carrol & Graf Publishers, Inc.: New York.

[10] Karlsson, Niklas, George Loewenstein, and Duane Seppi (2009). "The Ostrich Effect: Selective Attention to Information." Journal of Risk and Uncertainty 38: 95-115.

[11] Kopczuk, Wojciech and Joel Slemrod. (2005). "Denial of Death and Economic Behavior." Advances in Theoretical Economics 5(1): Article 5.

[12] Mullainathan, Sendhil and Andrei Shleifer. (2003). "The Market for News." American Economic Review 95(4): 1031-1053.

[13] Mynatt, C.R., M.E. Doherty, and R.D. Tweeney. (1977). "Confirmation Bias in a Simulated Research Environment: An Experimental Study of Scientific Influences." Quarterly Journal of Experimental Psychology 29: 85-95.

[14] Nickerson, Raymond S. (1998). "Confirmation Bias: A Ubiquitous Phenomenon in Many Guises." Review of General Psychology 2(2): 175-220.

[15] Rabin, Matthew, and Joel L. Schrag. (1999). "First Impressions Matter: A Model of Confirmatory Bias." Quarterly Journal of Economics 114(1): 37-82.

[16] Shaklee, H. and Baruch Fischhoff. (1982). "Strategies of Information Search in Causal Analysis." Memory and Cognition 10: 520-530.

[17] Snyder, M. and W.B. Swann. (1978). "Behavioral Confirmation in Social Interaction: From Social Perception to Social Reality." Journal of Experimental Social Psychology 14: 148-162.

[18] Stratton, Kathleen, Alicia Gable, Padma Shetty. (2001). Measles-Mumps-Rubella Vaccine and Autsim. Institute of Medicine, National Academy Press, Washington, D.C.

[19] Suen, Wing. (2004). "The Self-Perpetuation of Biased Beliefs." Economic Journal 114: 377-396.

[20] Swann, W.B., T. Giuliano, and D.M. Wegner. (1982). "Where Leading Questions Can Lead: The Power of Conjecture in Social Interaction." Journal of Personality and Social Psychology 42: 1025-1035.

[21] Wason, P.C. (1960). "On the Failure to Eliminate Hypotheses in a Conceptual Task." Quarterly Journal of Experimental Psychology 12: 129-140.

Figure 1 – Summary of Optimal Behavior

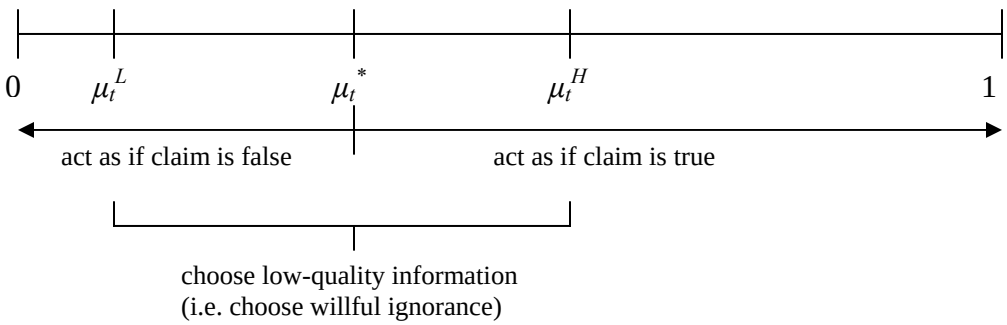

Figure 2a - Fraction of Population Choosing to Use Low-quality Information

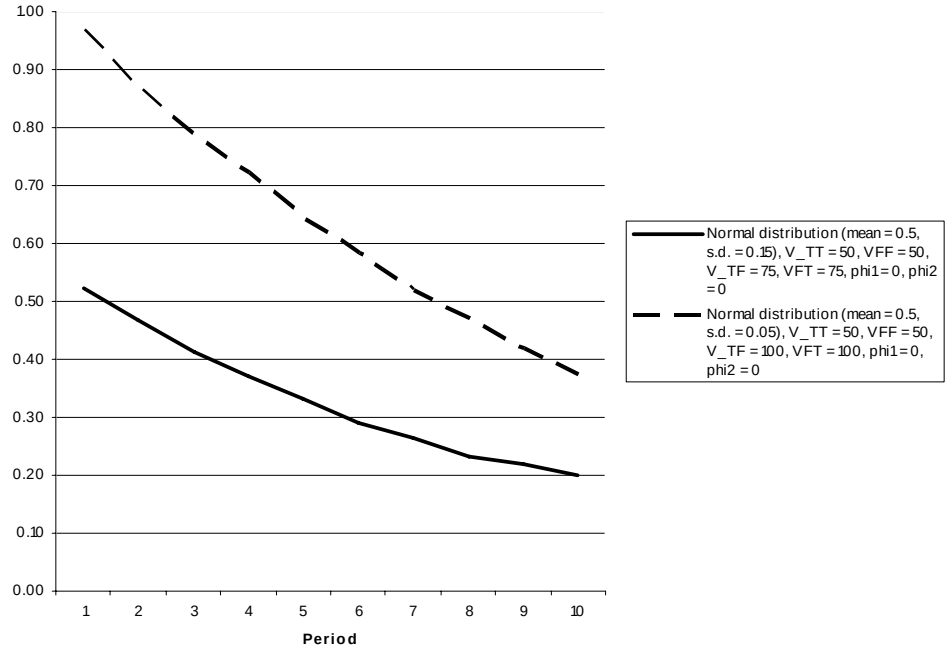

Figure 2b - Fraction of Population Acting in Opposition to the Truth

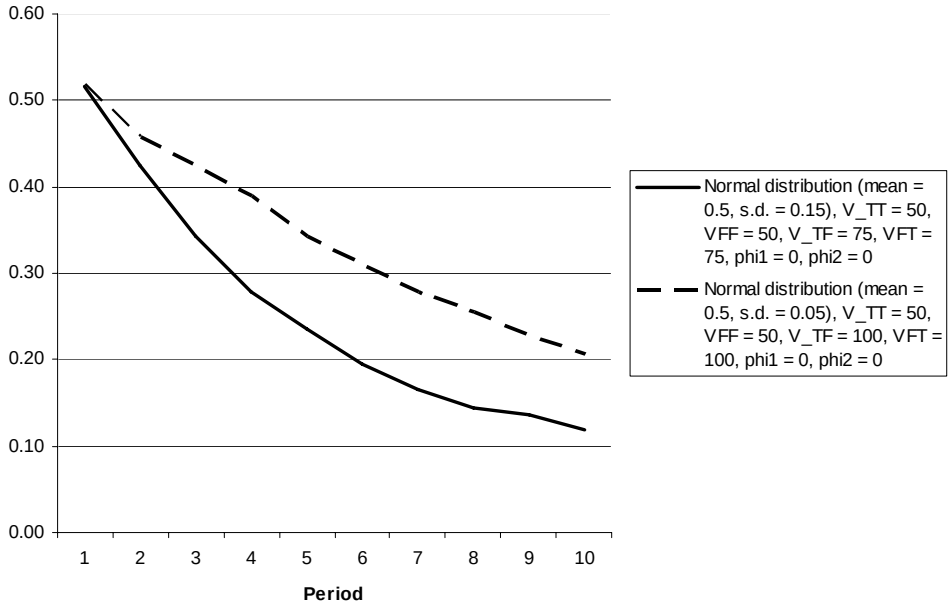

Figure 3a - Fraction of Population Choosing to Use Low-quality Information

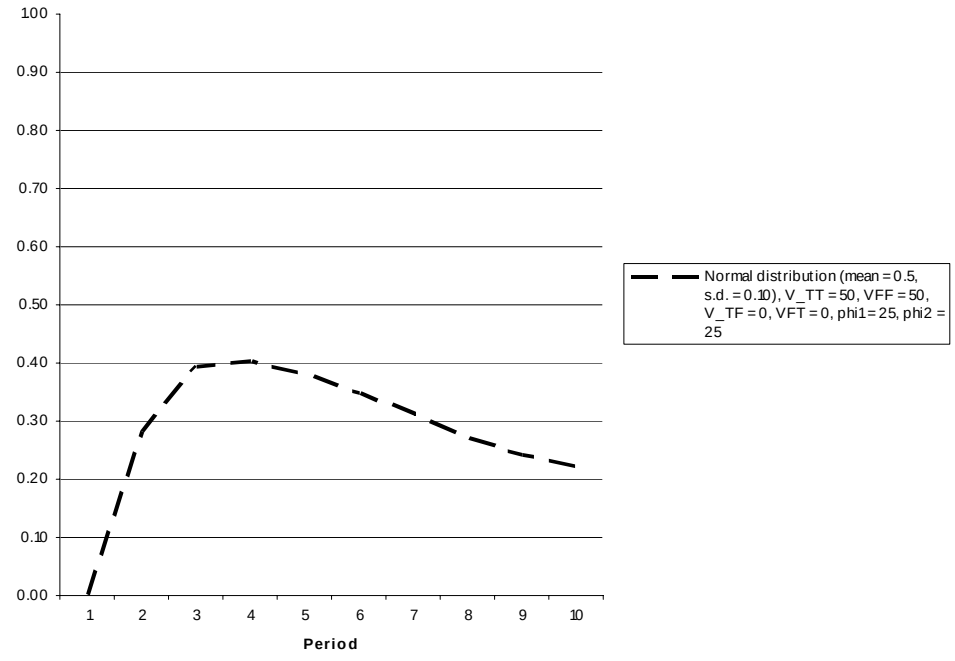

Figure 3b - Fraction of Population Acting in Opposition to the Truth