iSOLON.ORG
DEMOCRACY FOR THE INFORMATION AGE

WORKING PAPER

Automating Conflict-Of-Interest Reporting*

The FTC's new rules mandating blogger conflict-of-interest disclosure are based on antiquated technological assumptions.

By J.H. Snider, President
iSolon.org

March 10, 2010

Abstract

Countless laws across different local, state, and national governments mandate that service providers disclose material conflicts of interest to their customers. Countless ethics codes instituted by private entities such as trade associations and media outlets have similar but voluntary disclosure requirements.

Emerging web technologies, often categorized under the rubric of the "semantic web," offer a way to automate conflict-of-interest disclosure in new and powerful ways. To fully exploit these technologies, a well-structured vocabulary to describe conflicts of interest, a "bias ontology," would have to first be developed.

The efficiencies created through the use of such a technology would offer an opportunity to rethink the scope and methods of conflict-of-interest disclosure regulation. A promising place to start would be the FTC's conflict-of-interest disclosure laws, implemented December 1, 2009, for bloggers who review products.

As a complement to the current system of mandatory conflict-of-interest disclosure for legally recognized fiduciaries, a system of voluntary certification with significant penalties for false disclosure could be offered. In this way, service providers who wanted to make a credible claim that they are independent but are not currently covered by one of the government's mandatory conflict-of-interest laws could be offered an opportunity to do so.

* This working paper was released in conjunction with J.H. Snider's presentation, "Automating Conflict-of-Interest Reporting," at the U.S. Federal Trade Commission's workshop on "How Will Journalism Survive the Internet Age," Washington, DC, March 9-10, 2010. This paper stands alone but can best be appreciated when read in conjunction with that presentation. An early version of the argument contained in this paper can be found in J.H. Snider, "Automating Watchdog Reporting," *The Nieman Watchdog*," July 22, 2009. A discussion of the general conflict-of-interest problem in consumer protection can be found in J.H. Snider and Terra Ziporyn, *Future Shop: How New Technologies Will Change the Way We Shop and What We Buy*, an Authors Guild Backinprint.com Edition, 2008).

Automating Conflict-Of-Interest Reporting

The FTC's new rules mandating blogger conflict-of-interest disclosure are based on antiquated technological assumptions.

On December 1, 2009 the Federal Trade Commission (FTC) implemented a new rule that bloggers must disclose to their readers any cash or in-kind contributions they receive from companies whose products they review. Bloggers who fail to disclose conflicts of interest are subject to a penalty as high as $11,000 per non-disclosure.

The economic rationale for the new rule is that it reduces market inefficiency stemming from needless consumer confusion. According to this reasoning, American consumers face daunting product complexity in the marketplace—more than a million products, many with features that only an expert could evaluate. Bloggers who review products offer a promising new way to alleviate this confusion. But if bloggers can easily mislead readers about their degree of independence, then this promise will not be realized.

The FTC provided the following example to illustrate the impact of its rule on bloggers:

> A college student who has earned a reputation as a video game expert maintains a personal weblog or "blog" where he posts entries about his gaming experiences. Readers of his blog frequently seek his opinions about video game hardware and software. As it has done in the past, the manufacturer of a newly released video game system sends the student a free copy of the system and asks him to write about it on his blog. He tests the new gaming system and writes a favorable review. Because his review is disseminated via a form of consumer generated media in which his relationship to the advertiser is not inherently obvious, readers are unlikely to know that he has received the video game system free of charge in exchange for his review of the product, and given the value of the video game system, this fact likely would materially affect the credibility they attach to his endorsement. Accordingly, the blogger should clearly and conspicuously disclose that he received the gaming system free of charge. The manufacturer should advise him at the time it provides the gaming system that this connection should be disclosed, and it should have procedures in place to try to monitor his postings for compliance.

This rule is merely one of thousands of laws at the local, state, and national levels of government that require the disclosure of conflicts of interest by a wide variety of service providers, including doctors, lawyers, real estate brokers, financial advisors, radio broadcasters, and elected politicians. In addition, countless private entities, including trade associations and media outlets, have ethics codes with similar disclosure requirements.

This paper argues that current government approaches to mandating conflict-of-interest disclosure are technologically outdated and highly inefficient. It recommends that the FTC take a leadership role in updating conflict-of-interest disclosure in light of what new technology is rapidly making possible. Specifically, it recommends that the FTC undertake a broad rulemaking to rethink how information sources reveal and consumers access conflict-of-interest information. This rulemaking could cover not only bloggers but all service providers who are either required to or voluntarily claim to disclose conflict-of-interest information to their customers. The result could be a new and simplified consumer protection conflict-of-interest disclosure framework based, as described below, on semantic web technologies and a new voluntary certification option.

Bloggers who review products represent a good test case to modernize conflict-of-interest disclosure technology and policies. One reason is that bloggers, unlike many other service providers, rely exclusively on computers and the Internet to deliver their advice. Service providers, such as doctors, who offer their advice face-to-face, over-the-phone, or in handwritten communications, present a greater hurdle to modernizing conflict-of-interest disclosure. Another reason is that bloggers who review products are not engaging in political speech, which is more controversial to regulate because of First Amendment concerns.

The Promise of the Semantic Web

The Web is evolving from being exclusively human readable to being both human- and machine-readable. Today's web documents are easy for humans but not computers to extract meaning from. Semantic web technologies tag the data in documents, thus giving the data meaning (semantics) a computer can understand. Such tags on individual pieces of data are often called "metadata," and the logical linkages between the metadata "ontologies."

Computers may one day be as smart as humans in interpreting the meaning of documents without the help of explicit semantics such as metadata and ontologies. But until they are, adding such semantics to data will help computers process the data.

Note that although the term "semantic web" is often associated with a set of technologies developed by the World Wide Web Consortium, this paper uses that term in a generic sense to refer to any web-based technologies that use metadata and ontologies to enhance computer processing of data.

To understand the value of metadata, consider a news story that mentions the last name, "Branch." A human reading the story will probably immediately recognize from the context that the word "Branch" refers to a last name. But a computer processing the story may not be able to accurately infer whether "Branch" refers to a last name, first name, part of a tree, part of a river, a bank, religious sect, or some other type of object often described with the word "Branch." By adding metadata to "Branch" that identifies "Branch" as the last name of a person, this problem of machine interpretation can be solved.

When metadata are logically linked together to provide additional context, these larger contexts are called ontologies. Consider the product review ontology Google introduced in May 2009 as part of a rollout of what it calls "Rich Snippets." The product review ontology is comprised of the following types of metadata

1. The item being reviewed
2. The name of the item being reviewed
3. A numerical quality rating for the item being reviewed
4. The author of the review
5. The date that the item was reviewed
6. The body of the review

Google encourages websites with product reviews, such as Yelp!, TripAdvisor, and Epinions, to use the review ontology to tag the components of their reviews. It then picks up the tagged reviews when it crawls the Web to collect data for its search engine. The reviews are then used to improve consumer search results. Previously, when a consumer searched for a product, such as a restaurant, hotel, or model number, the search engine would display page links with snippets of text that described the product. Now, the search engines give higher priority to websites with reputations for useful product reviews, and the displayed snippets include starred product reviews, often with an average number of stars summarizing dozens of reviews for a particular product.

When combined with other data on the Web, such as Google Maps on an iPhone, the product review ontology is even more powerful, allowing consumers, for example, to easily display all pizza restaurants within a mile of their car. Figure 1 shows the search results for restaurants with three stars or more and within easy walking distance from the FTC.

In addition to product reviews, Google has endorsed a growing set of ontologies, including ontologies for "people," "products," and "organizations."

A critical feature of the new technology is that the user is given control over the interface. The consumer, for example, chooses exactly how the Google map appears and what types of restaurants are displayed. Yahoo!, which also crawls the web to aggregate product review data, adds a feature allowing consumers to choose a third party application to display the data.

Figure 1. Google's Product Review Ontology

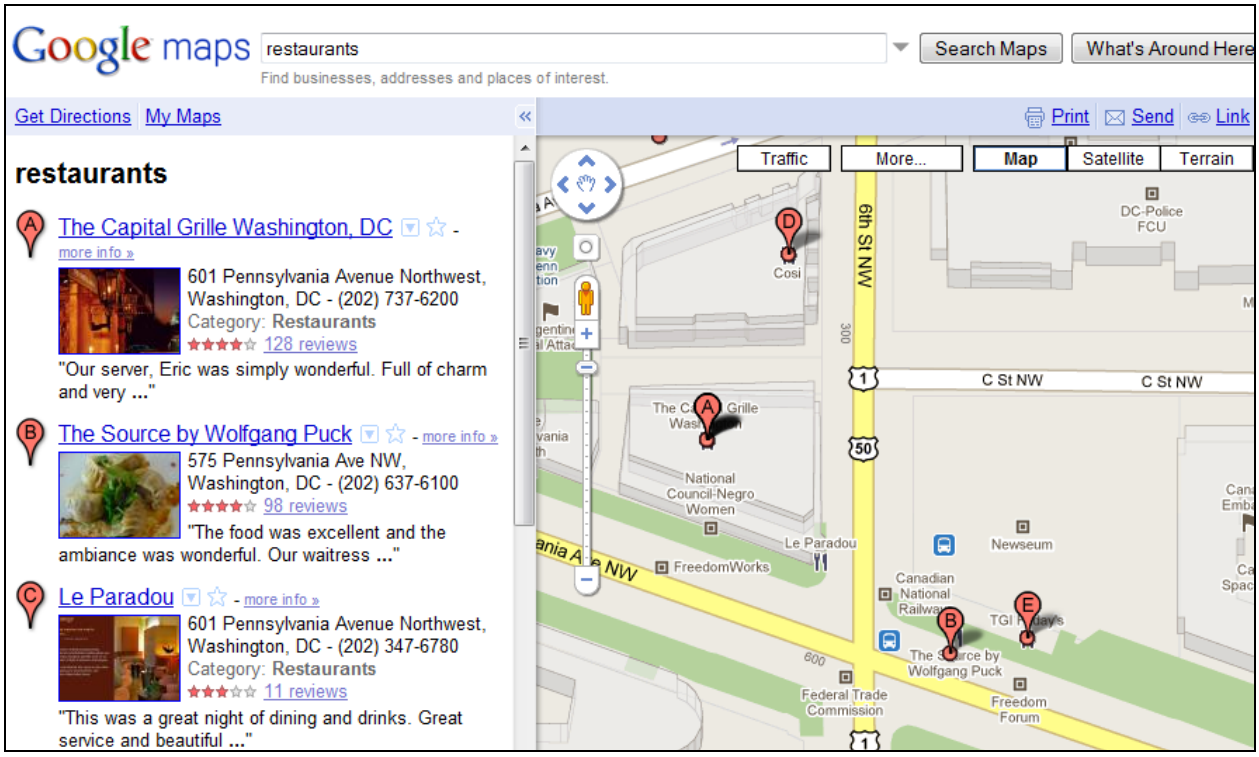

Metadata and ontologies need not be visible because they are intended for computer, not human, processing. In the case of a Web page, they can be hidden in the source code that generates the page. Figure 2 shows how the metadata for a product review is hidden in the source code. In this case, Google provides two different formats--microformats and rdfa--to tag the data.

One of the most widely adopted complex ontologies to date is XBRL, which describes business reporting data and has more than 14,000 metadata tags based on U.S. Generally Accepted Accounting Practices, also known as GAAP. The tags categorize potential line items on standard financial reports. This allows financial statements to be searched and displayed like relational databases. Versions of XBRL, all following the same basic logic of financial reporting, have been adopted by government financial regulators in more than a dozen countries, including the U.S., Netherlands, Israel, China, Australia, and Chile. In the U.S., the Securities & Exchange Commission (SEC) has required that public companies report their financial data in XBRL in a four year phased rollout beginning on June 15, 2009.

Figure 2. Metadata: For Machines, Not Humans, To Read

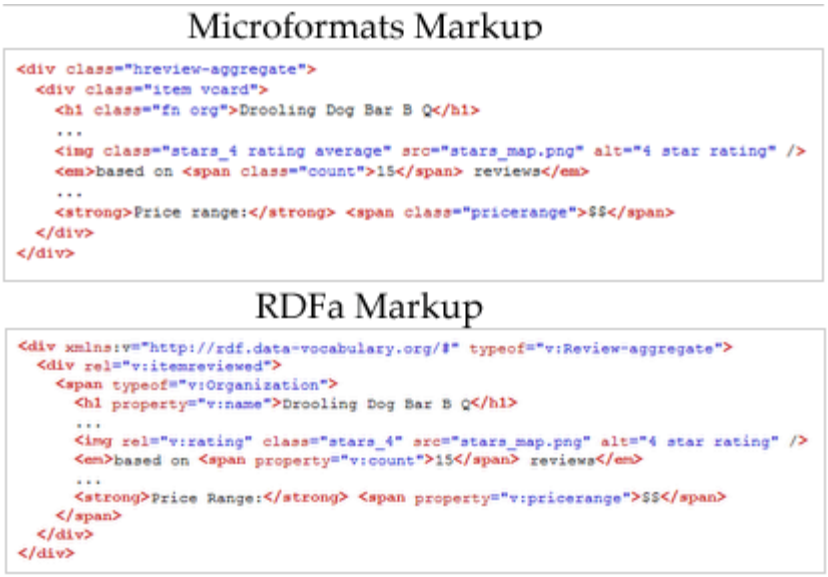

Bias Ontology

The display and access of conflict-of-interest information also needs to be rethought in light of semantic web technologies, starting with the development of an ontology, which I call a "bias ontology," to describe conflicts of interest.

The bias ontology is based on a subset of a well-developed academic theory known as principal-agent theory. In principal-agent theory, a principal delegates a task to an agent and a conflict of interest refers to any hidden source of compensation to the agent that would cause the agent not to act in the principal's interest. In the case of a blogger writing product reviews, the principal is the reader, the agent the blogger, and the conflict of interest a seller's hidden compensation to the blogger. Other examples of principal-agent relationships include patients delegating tasks to doctors, voters delegating tasks to representatives, shareholders delegating tasks to a board of directors, investors delegating tasks to financial advisors, home buyers delegating tasks to real estate brokers, and doctors delegating tasks to medical experts who publish their results in medical journals. For purposes of this paper, the terms information provider and information agent are interchangeable.

The two major components of a bias ontology are as follows:

1. Source of Agency Claim. This could be a government, agent, or another entity acting on behalf of the principal.

2. Agency Claim. This identifies the principal-agent relationship, such as a voter and elected official. The agency claim has five major subcomponents.

i. The principal. The principal can be an individual principal, such as a shareholder, or a collective principal, such as the class of all shareholders.

ii. Agent. The agent can be an individual agent, such as an elected official, or a collective agent, such as a legislature.

iii. Covered Interests. These describe an agent's income and assets.

iv. Covered Actions. These are the official actions, such as writing a blog, appointing a commissioner, or prescribing a drug, which an agent performs on behalf of a principal.

v. Default Settings Linking Covered Interests with Actions. These determine 1) the default covered interests that are automatically searched or manually entered when a covered action is taken, and 2) the way the resulting linkages are displayed.

A bias ontology would allow for increased automation of conflict-of-interest search. To investigate an agent's conflicts of interest, for example, the searcher would only need to identify the elements of an agency claim: the principal, agent, covered interests, and covered actions, with the search engine then automatically doing the thousands of queries necessary to look for potential conflicts of interest between the principal and agent.

In the case of doing a bias search on a government budget, for example, the increase of efficiency over current methods may be more than a millionfold. A budget for a large city or a small state may have more than a million transactions. Linking all those transactions to campaign contributors, gifts, and other potential conflicts of interest stored in separate databases now has to be done manually. This is so incredibly inefficient that it is rarely done. With a bias ontology, it could be done instantaneously. Figure 3 illustrates a linkage report between a covered action (a budget) and a covered interest (campaign contributions and gifts).

Another payoff from a bias ontology would be economies of scale in the applications market to make conflict-of-interest data useful. For example, one of the great benefits of XBRL becoming an international and widely used standard is that software companies around the world are now building sophisticated and affordable applications to both produce and consume XBRL data.

Figure 3. Elected Executive: Budget Conflict of Interest Alert

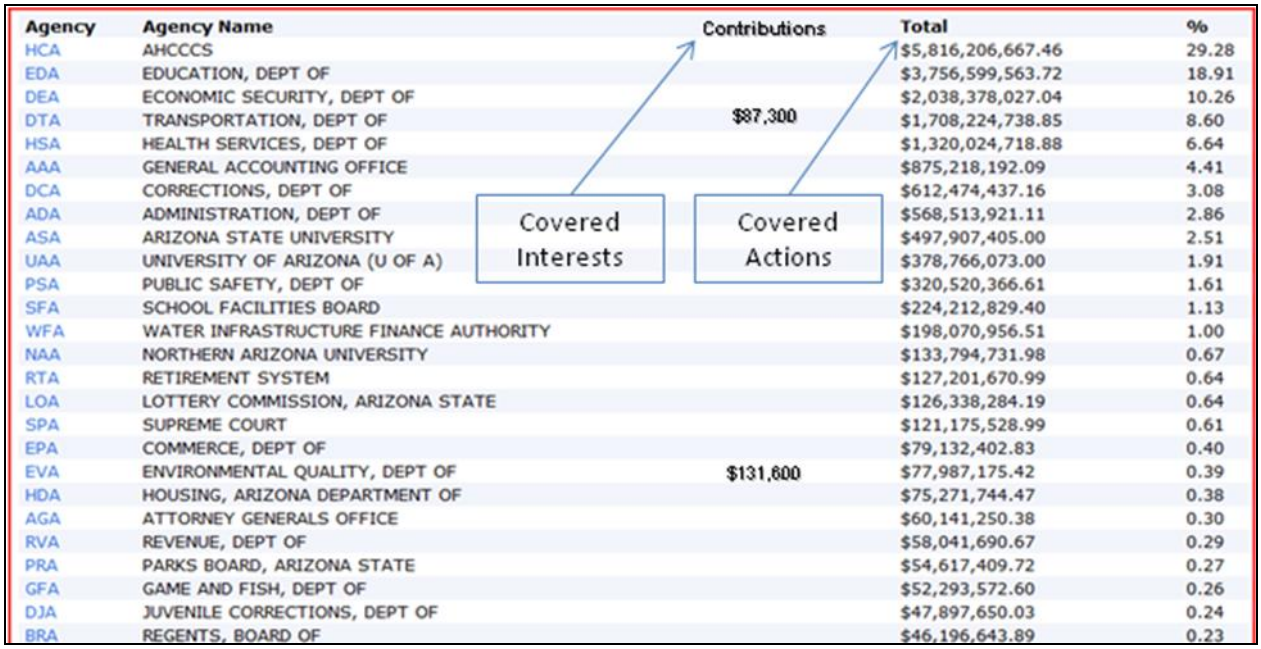

Agency	Agency Name	Contributions	Total	%
HCA	AHCCCS		$5,816,206,667.46	29.28
EDA	EDUCATION, DEPT OF		$3,756,599,563.72	18.91
DEA	ECONOMIC SECURITY, DEPT OF		$2,038,378,027.04	10.26
DTA	TRANSPORTATION, DEPT OF	$87,300	$1,708,224,738.85	8.60
HSA	HEALTH SERVICES, DEPT OF		$1,320,024,718.88	6.64
AAA	GENERAL ACCOUNTING OFFICE		$875,218,192.09	4.41
DCA	CORRECTIONS, DEPT OF		$612,474,437.16	3.08
ADA	ADMINISTRATION, DEPT OF		$568,513,921.11	2.86
ASA	ARIZONA STATE UNIVERSITY		$497,907,405.00	2.51
UAA	UNIVERSITY OF ARIZONA (U OF A)		$378,766,073.00	1.91
PSA	PUBLIC SAFETY, DEPT OF		$320,520,366.61	1.61
SFA	SCHOOL FACILITIES BOARD		$224,212,829.40	1.13
WFA	WATER INFRASTRUCTURE FINANCE AUTHORITY		$198,070,956.51	1.00
NAA	NORTHERN ARIZONA UNIVERSITY		$133,794,731.98	0.67
RTA	RETIREMENT SYSTEM		$127,201,670.99	0.64
LOA	LOTTERY COMMISSION, ARIZONA STATE		$126,338,284.19	0.64
SPA	SUPREME COURT		$121,175,528.99	0.61
EPA	COMMERCE, DEPT OF		$79,132,402.83	0.40
EVA	ENVIRONMENTAL QUALITY, DEPT OF	$131,600	$77,987,175.42	0.39
HDA	HOUSING, ARIZONA DEPARTMENT OF		$75,271,744.47	0.38
AGA	ATTORNEY GENERALS OFFICE		$60,141,250.38	0.30
RVA	REVENUE, DEPT OF		$58,041,690.67	0.29
PRA	PARKS BOARD, ARIZONA STATE		$54,617,409.72	0.27
GFA	GAME AND FISH, DEPT OF		$52,293,572.60	0.26
DJA	JUVENILE CORRECTIONS, DEPT OF		$47,897,650.03	0.24
BRA	REGENTS, BOARD OF		$46,196,643.89	0.23

In the case of conflict-of-interest data, consider the burden on consumers if every time they want to investigate a particular agent for a potential conflict of interest they must acquire and learn a different application. The conflict of interest data would then have much less value because it would be much more costly for consumers to use.

For it to be widely adopted, a bias ontology would need to be non-proprietary and flexible enough to cover all forms of principal-agent relationships. Only then could we expect a vibrant applications market to develop to both produce and consume the conflict-of-interest data. A precedent is the above mentioned adoption by government agencies of the non-proprietary and flexible business reporting ontology, XBRL. Like XBRL, a bias ontology could be an *extensible* ontology. This means that its basic deductive logic can be preserved even while its specifics, such as detailed lists of covered interests and covered actions, can be customized. Thus, for example, the FTC could develop custom fields to describe blogger conflicts of interest while being compatible with applications that cover many types of conflicts of interest, not just blogger conflicts of interest.

Figure 4 illustrates how a bias ontology might be used to highlight a potential conflict of interest within a blog product review. If the blog reader clicked on the highlighted text, it would list any monetary or in-kind compensation the seller of that product had given to the blogger. Note that the blog reader can be given complete control over how the conflict of interest is highlighted within the text. In this case, the potential conflict of interest is highlighted with a yellow marker in a red box around the reviewed product model reference. The reviewed product model reference could also have been highlighted with many other graphic techniques, including a footnote mark, different font, different color, or box

animation. Alternatively, the blogger's conflicts of interest could be aggregated in different ways and displayed in an alert, such as a sidebar, separate from the body of the text.

Figure 4. Blogger: Product Review Conflict of Interest Alert

Just as the SEC participated in an independent standards body, XBRL International, to develop XBRL for the financial reports of public companies, the FTC and other government agencies with consumer protection functions could participate in an independent standards body to develop a bias ontology. Once the core bias ontology is developed, government agencies and other organizations can then extend and customize it for their own applications.

A bias ontology cannot be designed and widely implemented in a day. As an analogy, Charlie Hoffman conceived of the business reporting ontology, later named XBRL, in 1998. But the SEC did not begin implementing it until 2009—eleven years later. Despite the long lead time and expense in developing XBRL, its government and private sector sponsors have

generally considered their investment worthwhile. Similarly, the benefits of a bias ontology could justify the necessary substantial and long term investment in it.

Certification vs. Mandatory Conflict-of-Interest Disclosure

The current FTC rules mandate conflict-of-interest disclosure. As an alternative, the FTC and other government agencies could create a new type of voluntary certification program for bloggers and other information providers who want to be able to make a credible claim that they are independent.

The certification would designate a new type of transparent organizational entity, called an "independent agent." Independent agent certification would be strictly voluntary but would entail significant penalties if the terms of certification were violated. Bloggers and other information providers would evaluate the benefits and costs of a legally binding conflict-of-interest disclosure requirement before deciding to become an independent agent. The agent's claim to be independent would be written in the bias ontology.

An analogous certification process is incorporation. The government doesn't mandate the corporate form of organization in order to conduct business. But most companies decide to incur the high upfront and ongoing costs of incorporation because of the many benefits that come from such certification. For example, incorporation exempts shareholders from personal liability if a company goes bankrupt. This greatly increases the incentives for people to invest their capital in companies, which is good for society.

A blogger or any other information provider's advantage in becoming a certified independent agent would be its enhanced ability to credibly claim that it is independent. The enhanced credibility would arise because the agent would have a higher probability of incurring a significant penalty for not disclosing all material conflicts of interest in a way consumers could easily discover and evaluate.

Too often, mandatory conflict-of-interest rules have been designed and implemented on an ad hoc basis in response to narrowly defined scandals and special interest pressures rather than market needs. The resulting gaps in effective disclosure rules create highly distorted winner-take-all information markets that place a premium on promoting rather than creating innovative information products and leave American consumers information impoverished. By offering *all* information providers a cost-effective option to credibly demonstrate their independence, voluntary certification would help fill in many of these gaps.

Voluntary certification would also minimize First Amendment concerns based on government regulating information sources. From a First Amendment standpoint, it would be heinous to mandate that all news outlets disclose conflicts of interest. But if a particular news outlet, such as a news outlet specializing in consumer finance, believes that its claims

of independence won't be credible with its audience, it may welcome the credibility that government enforcement can bring.

Independent agent certification would most benefit small information providers, such as bloggers, who currently find it prohibitively expensive to establish consumer trust. Large, well-established and trusted information provider brands, such as the *New York Times*, *Wall Street Journal*, or *Washington Post*, would presumably have little to gain from such certification and would thus not seek it. However, to implement their ethics policies more efficiently, effectively, and credibly, they still may want to use a bias ontology for conflict-of-interest disclosure.

Companies with a narrow business line focused on consumer information would also be more likely to benefit from independent agent certification than large, diversified companies with many intrinsic conflicts of interests. The more conflicts of interest a company has, the less it would presumably want to legally bind itself to disclosing them.

Similarly, information providers that cover politics may, on First Amendment grounds, be more resistant to any government involvement in conflict-of-interest disclosure than information providers covering commercial products.

The FTC's new rules, by excluding mainstream media and focusing on bloggers writing product reviews, implicitly already make such distinctions. But by making the rules mandatory on bloggers and not other information providers, including consumer magazines and others that specialize in product reviews, the FTC opens itself to charges that its rules are arbitrary and unfair.

Old vs. New Blogger Disclosure Scenarios

Let's now compare two scenarios: the FTC's current plan for blogger conflict-of-interest disclosure and the plan proposed here based on an open, standardized bias ontology and voluntary certification.

Scenario #1: The FTC's Current Rules. Sally writes a blog, Baby Products Review, that reviews baby products. Toddler, Inc. gives Sally a free copy of Baby Product Y. Sally is legally required to disclose in a review of Baby Product Y that she received a free sample. She controls how this information is displayed, providing it is displayed prominently.

Toddler, Inc. is required to monitor Sally's blog to ensure that the gift has been disclosed.

Jane, Sally's reader, has minimal control over whether or how she wants to see this information.

The FTC must manually monitor Toddler, Inc. and Sally's blog to enforce compliance. After collecting this data, the FTC has great discretionary power to assess whether Toddler, Inc.

and Sally have indeed complied with its laws. The resulting monitoring and compliance cost is so high that it is only practical to enforce egregious violations resulting in significant financial penalties.

Scenario #2: The Proposed FTC Rules. Sally writes a blog that reviews baby products. She wants to win her readers' trust in her objectivity, so she registers her blog, Baby Products Review, as an independent agent. When an agent for Toddler, Inc. gives Sally a free copy of Baby Product Y, the agent and Sally are each legally obliged to publish the gift information to the Web in a bias ontology, which is then collected into a centralized database. The double entry requirement is akin to the IRS mandating that both employers and employees report the same income separately; each entry serves as a check on the other and makes collusion to hide information more costly.

The burden on Toddler, Inc.'s agent is now greatly reduced because it now only needs to worry about its own disclosure of a potential conflict of interest; it doesn't have to monitor Sally's blog for inadequate disclosure. When Sally writes her review of Product Y, she must include a unique identifier that can be linked back to the centralized conflict-of-interest database. And if Sally fails to do this, others, including competitors, will be able to automatically consult the centralized database and report the omission to the FTC.

Jane is now given complete control over how the conflict-of-interest information is displayed when she reads Sally's blog. No longer does Sally determine the place and style of her conflict-of-interest disclosure and no longer does the FTC have to write detailed rules describing what "prominent" disclosure means. If Sally wants a summary of the conflicts of interest at the beginning, side, or end of the review, she can so choose. If she wants the conflicts of interest disclosed in the text where the product is mentioned (e.g., with a particular color, font, underline, animation, or footnote), she can do so. If she wants to determine the materiality of the disclosed conflicts (e.g., "I only care about product gifts worth more than $25"), she can do so. And if she wants to customize how conflicts are disclosed for particular blogs or all the blogs she visits, she can do so by setting the preference to her conflict-of-interest viewer within her Web browser.

With this new system, we now have new types of independent, technology-based conflict-of-interest players: 1) aggregators of conflict-of-interest data, and 2) application developers that process and display conflict-of-interest-data. A single entity may also provide both functions. The aggregators, such as Google or Bing, collect all the conflict of interest data posted to the Web into a consolidated database. They do this automatically and behind-the-scenes. The application developers process and display the conflict-of-interest data on behalf of readers using a simple and convenient interface such as a browser plug-in.

In Jane's case, she chooses her own conflict-of-interest-disclosure plug-in and sets its preferences once, then uses it for all her web browsing. When Jane opens the web page of a

blogger registered as an independent agent, the plug-in works behind the scenes on behalf of Jane, making countless automated queries to check for all manner of conflicts of interest regarding product endorsements. The plug-in then displays the material conflicts of interest based on Jane's pre-established preferences.

As for the FTC, it delegates a much greater share of its compliance monitoring costs to Jane, competitors to Toddler, Inc., and competitors to Sally's Baby Products Review. This is possible because these entities now have greatly reduced costs of monitoring and reporting on the independence of Sally's Baby Products Review. The FTC also writes disclosure laws that are more precise (because they need to be precisely specified to be machine readable) and limited (because otherwise controversial human interface design issues are left to Jane).

The FTC also has a larger set of sanctions if Sally fails to disclose her conflicts of interest. In addition to its current menu of monetary penalties for gross abuses, it can provide a violation notice, including the number of violations and links to their details, that consumers could automatically see when reading Sally's blog. It can also withdraw Sally's status as an independent agent either temporarily or permanently. If Sally had built up a lot of goodwill in her blog, such a revocation of independent agent certification could be very costly for her.

Mandatory Certification for Other Consumer Markets

For licensed occupations that have the legal obligations of a fiduciary, including a duty of loyalty to their principals, the use of a bias ontology could be made mandatory. Examples of such occupations include financial advisors, real estate brokers, and doctors. Of course, the disclosure would depend on the type of covered action the agent undertook and thus be different than for a blogger who writes product reviews. For example, doctors who write patient prescriptions could be required to write prescriptions using a bias ontology and then send an electronic copy to the patient so, like a reader detecting a conflict of interest in a blogger's product review, the patient could automatically detect a conflict of interest in the doctor's prescription.

In the many cases where fiduciaries already are legally required to disclose conflicts of interests, the laws are scattered in a hodgepodge of obscure documents and administered by a hodgepodge of obscure government agencies. A bias ontology would translate all these laws into a standardized, structured, and machine readable format.

Private entities may want to make use of a bias ontology as a condition of membership, funding, publication, or some other privilege. For example, trade associations may want to restrict membership to those who comply with their ethics guidelines by incorporating a bias ontology in business communications; foundations may want to restrict grants to nonprofits that incorporate a bias ontology in their nonprofit communications; and academic journals may want to require authors to submit financial conflict-of-interest information in a bias ontology.

Caveats

Many conflicts of interests are subtle and require great judgment to evaluate. Even a well executed bias ontology can only point to possible conflicts of interest. Humans would still be required to evaluate their significance.

A bias ontology can only generate correlations between covered interests and covered actions; it cannot prove causality. Statistical techniques would need to be applied to the data to infer causality.

Even a well executed bias ontology can generate many false positives or false negatives, depending in part on whether consumers set a high or low conflict-of-interest disclosure hurdle. Like adjusting the cookie settings on their computers, consumers would have to tinker with their conflict-of-interest disclosure settings to avoid getting either too many false positives or negatives.

Judgment calls would still have to be made on what constitutes a material conflict of interest subject to disclosure. Given that modern information technology can automate much of the disclosure process for sources, government agencies would presumably be more willing to err on the side of too much disclosure. For example, instead of setting a gift disclosure bar at $25, it might set it at $5 and let consumers decide whether they consider gifts worth as low as $5 to be material. Alternatively, the information source could specify what it defines to be a material conflict of interest. The consumer could then avoid information sources that set the bar of materiality too low.

Conclusion

Technology has long been a driver of regulatory policy. For example, implementing the FTC's Do Not Call Registry prior to the advent of the personal computer and Internet age would have been much more costly for households, callers, and the FTC, thus probably rendering the Registry an impractical innovation.

Similarly, a bias ontology technology offers a whole new menu of regulatory possibilities. To bring these possibilities to fruition, the following two general types of steps should be taken.

First, develop a bias ontology as a joint effort of the government and private sector. This joint effort could include setting up a standards body to design the bias ontology. The FTC's conflict of interest rules for bloggers writing product reviews could be a good test case to implement the bias ontology.

Second, create a system of certification for agents who are seeking a more efficient mechanism to make a credible commitment that they are independent. Offering certification

to bloggers who review products could be a good test of the appeal and effectiveness of voluntary certification.

The vision of automating conflict-of-interest reporting outlined here is just the tip of the iceberg of how the FTC, Congress, the White House, and state and local government need to rethink consumer protection in light of emerging semantic web technologies. America's Rube Goldberg system of consumer protection needs to be rethought from the ground up based on semantic web technology. Executed well, a bias ontology could point the way.