

The Informed Press Favored the Policy Analysis Market

Robin Hanson*
Department of Economics
George Mason University†

August 8, 2005
(First Version March 18, 2005)

Abstract

The Policy Analysis Market (PAM), otherwise known (inaccurately) as “terrorism futures,” burst into public view in a firestorm of condemnation on July 29, 2003, and was canceled within one day. We look at the impression given of PAM by five hundred media articles, and how that impression varied with eleven indicators of article information quality: publication date, citing someone with firsthand knowledge, article length, a news or an editorial style, author anonymity, and the awards, circulation, frequency, and topic specialties of the periodical. All eleven indicators individually predict more favorable impressions of PAM. In a multiple regression model, seven of them remain clearly significant. This model predicts that a fifty word news article in an award-winning widely-read science and business publication a month later that mentioned an insider would give a solidly favorable impression of PAM, as would a similar three hundred word article today in a general publication.

*For their comments, I thank Bryan Caplan, Rebecca Jarmas, John Lott, Christian Masse, Ute Shaw, and Justin Wolfers. For research assistance, I thank Nate Russell, Bonnie Hanson, and “Chris. F. Masse.” I thank the Center for Study of Public Choice, the Mercatus Center, and Microsoft for financial support.

†rhanson@gmu.edu <http://hanson.gmu.edu> 703-993-2326 FAX: 703-993-2323 MSN 1D3, Carow Hall, Fairfax VA 22030

Introduction

On July 28, 2003, two U.S. Senators announced that a Pentagon betting market on terror attacks was about to open, which terrorists could abuse. Amid widespread condemnation, this project was canceled the next day, and two days later its widely reviled supervisor John Poindexter resigned.

The Policy Analysis Market (PAM), however, was a U.S. military research project designed to test the ability of speculative markets to forecast overall geopolitical trends, not terror attack details (Polk, Hanson, Ledyard, & Ishikida, 2003). This two-year-old million-dollar research project began long before Poindexter joined its funder, the Defense Advanced Research Projects Agency (DARPA), and was five months from public trading of under one hundred dollar bets. PAM traders could have, for example, bet on the chance of high levels of civil unrest in Saudi Arabia in the fourth quarter of 2004, conditional on the US moving its troops out of there two quarters earlier. By comparing estimates based on different assumptions, PAM could have advised us on the effect of various US Mideast policies.

The Democratic Senators' complaint seems to have been deliberately misleading, in order to embarrass Bush via his association with the previously reviled and now newly tainted Poindexter. The PAM webpages consisted of text describing PAM, shown over faint backgrounds of sample PAM interface screens. In addition to large sections on geopolitical trends, two of these sample screen contained a small ($< 2\%$) miscellaneous section, with short phrases about a possible Arafat assassination, North Korean missile attack, and the king of Jordan being overthrown¹

The Senators' complaint was timed to appear when the DARPA public relations person was out of town and out of reach. The political and public reaction was immediate, strong, and negative². The moral outrage factor forced the Bush administration to make a fast decision to embrace or reject PAM, and the timing forced this decision to be based on very little information about PAM.

As a participant in the PAM project, I followed the media coverage with great interest, and it seemed to me that the coverage became more positive as it became better informed. In this paper, we test this hypothesis, and study what indicates articles that are more favorable impression of PAM. We had four people each rate over five hundred media articles that mention PAM, on a seven point scale in terms of how favorable or unfavorable an impression of PAM they give. We then looked at how the average favorability rating of an article varied with eleven indicators of article information levels, along with six other relevant variables.

These information indicators were: elapsed time, article length, whether the publication specialized in science or finance, the period between publications, whether the periodical had many readers or many journalism awards, whether the article was written as news reporting

¹Also, on May 20, 2003, a DARPA report to Congress described the larger FutureMAP program using the example "Will terrorists attack Israel with bioweapons in the next year?"

²Three web polls are available: *Excite.com* on July 30, 2003 was 88% against PAM, *AJC.com* undated gave 49% "bad idea" to 37% "good idea," out of 591, and *TheWBALChannel.com* undated gave 68% "absurd" and "wrong" to 28% "I would invest" out of 1705.

or as an editorial, whether it mentioned any insider who was directly involved with PAM, and whether the author was identified. The other control variables were author gender, publication political leaning, and the media type, including whether it was print, broadcast, or web only publication.

Taken individually, each of the eleven indicators of more informed articles favored PAM, suggesting that in general more informed articles gave a more favorable impression of PAM. When combined into a multiple regression, seven of these indicators remained clearly significant, and another two were marginally significant. The model predicts that a fifteen hundred word editorial on July 30, 2003 (i.e., the two days later) that mentioned no insiders would give a solidly unfavorable PAM impression. But a sixty word news article in an award-winning widely-read science and business publication a month later that mentioned an insider would give a solidly favorable impression, as would a similar thousand word article today in a general topic publication.

If the strong correlation among information indicators found in this study continues to hold on other topics, the method used in this study might well allow us to identify the more informed sides in future controversies as they happen.

Previous Research

Journalism researchers have long studied reporting errors and their indicators in newspapers, magazines, and television (Singletary, 1980), finding that about half of straight news stories contain an error. Researchers have found higher accuracy in good relative to bad news, in articles whose sources are press releases and court records relative to police, letters, and other articles, in ordinary news relative to reporting on science or social issues, in longer relative to shorter quotes (Burriss, 1985), in prestigious relative to high circulation newspapers (Lacy, Fico, & Simon, 1991), in unanticipated relative to anticipated news (Berry, 1967), in face-to-face relative to telephone interviews, by reporters with weak relative to strong opinions on the topic, in bylined relative to unsigned articles (Singer, 1990), by topic-specialized relative to general reporters, by reporters with more relative to less neurotic personality types, and by passively observing relative to actively interviewing reporters (Scanlon, 1972). On the other hand, no relation has been found between accuracy and degree of personality introversion, the type of source used (e.g., government or business), the frequency with which a person is cited, or the amount of contact between a reporter and a source (Maier, 2002; Burriss, 1985). There are conflicting reports on the relative accuracy of different media types; some analyzes prefer weekly magazines (Moore & Singletary, 1985; Dunwoody & Scott, 1982), some prefer newspapers (Hanson & Wearden, 2004), and some prefer television and wire services (Singer, 1990).

Librarians have long distributed guidelines to help their patrons judge the relative reliability of differing publications, such as web sites (Pask, Kramer, & Mandernack, 1993; Furno, Bolton, & Heim, 2002). For example, these guidelines prefer publications in a neutral news or analysis style relative to more opinionated styles, which cite more references (especially original sources), that have a more recent publication date, and that are written by a named

author with high educational credentials who is not selling something. While there is little research to explicitly validate these guidelines, they are certainly reasonable candidates to consider.

Economists have recently begun to study political media bias (Groseclose & Milyo, 2004; Lott & Hassett, 2004), but have so far paid little attention to media information-quality indicators.

There have been two previous published analyzes of the PAM press. Two weeks after the Senators' complaint, an informal media analysis concluded that:

The resulting uproar and subsequent termination of the terrorism futures market served as an excellent example of the power of the media and the impact of public opinion ... Media coverage was replete with indignation that the Pentagon would even consider creating and funding what several called a morbid betting parlor ... Nearly two-thirds of the stories either defended the program or gave a more detailed explanation of what the perceived usefulness of the program was meant to be. ... Approximately half of the analyzed reporting focused on PAM being run by retired Admiral John Poindexter, ... While thinking outside the box is normally seen as a good thing, there appeared to a very broad consensus that in this case Poindexter and other officials went way too far (Clifton, 2003).

Another team reported on a statistical analysis of the favorability of PAM press and a single information indicator. Joshua Tucker and Adam Meirowitz made a three point classification (opposed, neutral, in favor) of 115 articles from July and August 2003. They also classified them as news (48 articles) or opinion (67 articles). They found that opinion articles were mostly opposed, while news articles were mostly neutral. Of the 9 non-neutral news articles, 7 were opposed, while of the 63 non-neutral opinion articles, 49 were opposed. News articles were on average more favorable, but one might interpret this as because news favors a neutral stance to all topics. After all, they were both 78% opposed when they had a non-neutral position (Tucker & Meirowitz, 2004).

The Data

We have collected a publicly-accessible Policy Analysis Market Archive³, which includes copies of all media articles that we have found that mention the Policy Analysis Market in any recognizable way. (Our search included looking in all standard media databases for "terrorism futures," "terror betting," "policy analysis market," and "FutureMAP")

We recruited four people to read or listen to all 555 (non-blog) media articles that we could find and make available online.⁴ They happened to be a female Democrat, age 73, a male Republican, age 57, and two political independents, a male, age 23, and a female, age

³Available at <http://hanson.gmu.edu/policyanalysismarket.html>.

⁴A single Japanese language article was rated instead by person of Japanese descent.

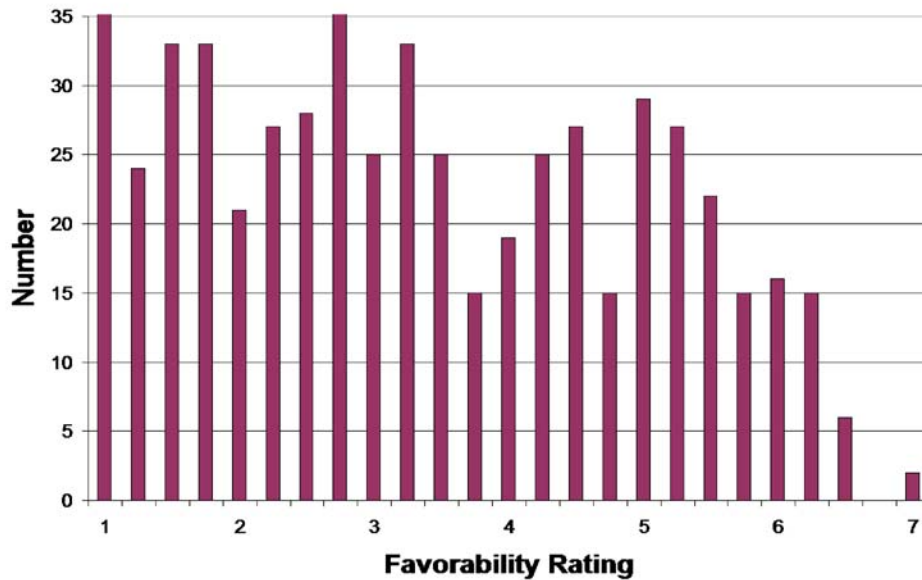


Figure 1: Rating Distribution

26.⁵ They were asked to rate the articles on a seven point scale regarding how favorable or unfavorable an impression of PAM an ordinary person might get from that article. That scale was: (1) very unfavorable, (2) unfavorable, (3) somewhat unfavorable, (4) neutral, (5) somewhat favorable, (6) favorable, (7) very favorable.

The pairwise correlation between the reader’s ratings ranged from 0.726 to 0.759, and the average of the ratings became each article’s favorability rating. The average favorability in the sample was 3.40, or a bit closer to somewhat unfavorable than to neutral, and the standard deviation was 1.59. Figure 1 shows the full distribution of average ratings.

In addition to the article rating, we collected eleven variables that can arguably serve as indications of how informed an author was about PAM when the article was written. These are not the only indicators that one could collect, but they were among the cheapest for us to collect. These indicators can be categorized according to the strength of our prior expectations about their relevance.

We expected the strongest indicators of the information level of an article to be its date and length, whether it cited sources with first-hand knowledge, and the topical expertise of the publication it appeared in.

Information should have become more easily available with time, in part because more previous articles were available for reference. We calculated the number of days from the Senators’ complaint on July 28, 2003 to the article’s appearance. If the big complaint misled the media for a time, articles should have been more favorable both before the misleading

⁵All four were Caucasian Christians.

complaint and as time passed after that complaint.

One expects that longer articles are on average researched in more depth, at least regarding the main topic of the article. It might be preferable to measure only the length of the sections on topics “related” to PAM, but this can require many subjective judgments. Thus when there were clear article boundaries, the length of the entire article in words was calculated. This was done exactly for text articles, and estimated for audio articles. For radio or TV transcripts that ranged widely in topic without clear article boundaries, only the section dealing with PAM was calculated. For books, the length in words of the entire book was estimated from the page count.⁶

Since PAM was a technology research project involving markets, publications which specialize in technology or markets might have plausibly had more relevant expertise. We coded whether a publication specialized in science or technology, and also coded whether it specialized in business or finance. Articles in the business section of a general newspaper were also counted as business articles.

Someone trying to learn the truth about PAM might have talked to one of the insiders involved in the project, such as Michael Foster, the DARPA manager of the FutureMAP program which funded PAM, and the main PAM project participants, i.e., Robin Hanson, John Ledyard, and Charles Polk.⁷ We recorded whether each article mentioned any of these names of PAM insiders. (To avoid biasing the results, this study excludes all writings by PAM insiders.) Of course it could also be that someone who for other reasons wants to present a favorable picture will be more likely to want to quote the likely more favorable project insiders.

In addition to these strong indicators, we had some moderate information indicators, where our expectations had a clear sign, but were not as strong. Such indicators included whether an article was in a news or an editorial style, whether it was in a high circulation or award winning publication, how frequently the publication appeared, and whether the author was identified.

News reporters are supposedly held to higher standards of analysis and objectivity than columnists and opinion writers. For each article we recorded whether or not it was clearly in a news reporting style (i.e., neutral and descriptive, versus opinionated and interpretive), and so likely to be held to such a higher standard. Lead editorials, op’eds, interviews, columns, reviews, academic papers, and reports were not coded as news. Since lead editorials, endorsed by an editor or editorial board, are expected to be especially opinionated, we also coded whether an article was such an editorial. Most articles were neither clearly news nor lead editorials.

Reporters from more prestigious publications may be held to higher standards of analysis and objectivity. We used two indicators of prestige, the publication’s circulation, and how many prestigious news awards it has received. We coded whether an article appeared in a

⁶Length is an especially bad approximation for books, since most books only had a paragraph or so on PAM. We thus controlled for whether something was a book in our data analysis.

⁷John Poindexter was not an insider, as he had little knowledge of project details.

print publication with a circulation of over 900,000 readers.⁸ We also coded whether a print publication was one of the top ten historical finalists of the Pulitzer Prize, The National Magazine Awards, or one of the top three of the British Press Awards.⁹

One might expect more depth of analysis from periodicals with longer periods between publication, if reporters who have more time before their deadline can put more thought into their words. The publication period in days was used when available. Books and book chapters were set as having a publication period of 180 days, while reports were set to 90 days.

When articles do not have a credited author, the true author might be less concerned about how the article will effect his reputation. For each article we coded whether it was anonymous, or had specific credited authors.

In addition to these strong and weak information indicators, we also had some variables that seemed to be plausibly relevant, but not necessarily as information indicators. That is, for these variables we had at most only very weak expectations about which side would be more informed. These variables included author gender, and the media type, frequency, and political leaning of the periodical.

There have been many conflicting claims about which media types are more informative. We coded whether an article was from a wire service, whether it appeared in a broadcast media of TV and radio, and whether it appears only on the web or email. All others are classified as “print” articles.

Gender is a relevant variable in many contexts, and some studies have found men to be on average better informed than women (Caplan, 2001; Kraus, Malmfors, & Slovic, 1995; Delli-Carpini & Keeter, 1997).¹⁰ For articles with credited authors, we coded whether that author was male or female. For articles with multiple authors we averaged their gender, and for anonymous articles we used the average gender of the rest of the sample.

Finally, we made an effort to control for political leanings. A recent clever analysis, using citations of think tanks, has allowed a standard rating of congressman liberalism, the ADA (Americans for Democratic Action) score, to be projected onto twenty popular news sources (Groseclose & Milyo, 2004). The mean Senator ADA score is 40, the mean House of Representatives ADA score is 44.5, and the twenty media sources studied had a mean of 62.6 and a standard deviation of 11.0. We found that 101 of our 555 articles could be matched

⁸In this sample, these publications were: The Mail, The Mirror, USA Today, Wall Street Journal, New York Times, Los Angeles Times, The Telegraph, The Express, Time, Newsweek, US News & World Report, Kiplinger’s, Business Week, and New Yorker.

⁹The top ten Pulitzer historical finalists are New York Times (140), Washington Post (92), Los Angeles Times (86), Philadelphia Inquirer (63), Wall Street Journal (53), Boston Globe (48), Chicago Tribune (47), St. Louis Post-Dispatch (26), Seattle Times (20), and Chicago Sun-Times (14). The top ten National Magazine finalists are: New Yorker (150), Esquire (71), Newsweek (44), Time (39), Business Week (35), Fortune (17), US News & World Report (15), New Republic (15), Forbes (11), and tied for tenth, Wired and IEEE Spectrum (10). The top three British Press finalists are the Times (49), Telegraph (37), and Guardian (37).

¹⁰I was prodded to add a gender variable by reading a female-authored opinion article which suggested that women are more careful when writing opinion articles (Smardz, 2005).

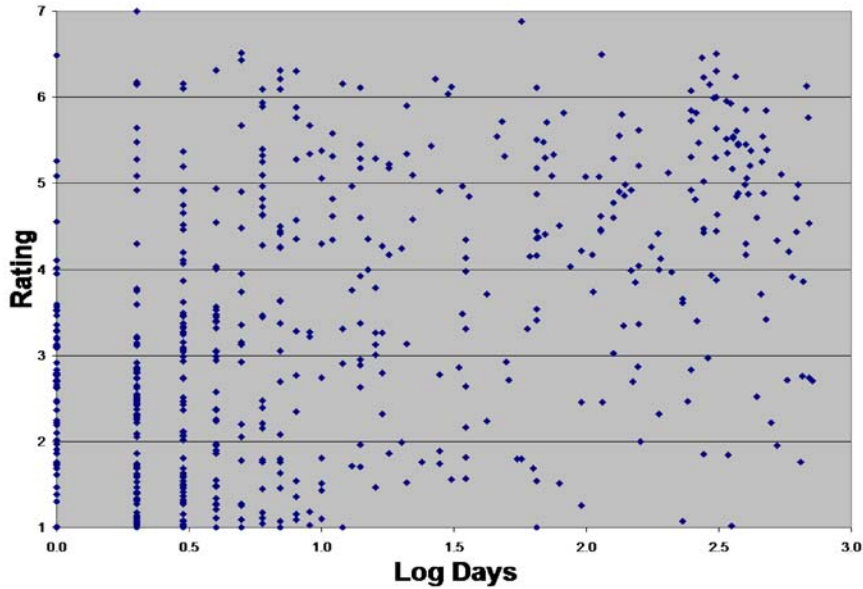


Figure 2: Favorability Rating Vs. Date of Article

with one of these twenty rated media sources.¹¹ The remaining articles were given the mean media ADA score of 62.6.

Results

Table 1 shows how average favorability depended individually on these various indicators of article information. For each category, table 1 shows the number of articles and their mean favorability relative to the average of the entire set of 555 articles. Positive entries flag categories that are more favorable than average; negative entries are less favorable than average.

Note that average opinion as a function of time fluctuated up and down, as it should if it reflected a random-walk-like information revelation process.¹² Figure 2 shows more detail about dependence on the (base ten) logarithm of an article’s days from the Senators’ complaint, while figure 3 shows more detail about dependence on the (base ten) logarithm of article length.¹³

Surprisingly, in table 1 *all* eleven of the proposed information indicators, both strong

¹¹For three TV networks the study distinguished different scores for different shows. In those cases we averaged them to give a single network ADA score.

¹²Note also that the time trend is not just a matter of a jump after the project was canceled on day one.

¹³In figures 2 and 3, a random number in $[-0.125, 0.125]$ was added to each rating values to allow the data points to be better distinguished visually.

<i>Days After 7/28/2003</i>	<i>Count</i>	<i>Deviation</i>	<i>Period</i>	<i>Count</i>	<i>Deviation</i>
Before	7	1.00	Daily	404	-0.27
1	57	-0.64	Weekly	67	0.79
2	80	-0.87	Bi-Monthly	6	1.10
3	61	-0.63	Monthly	37	0.74
4	38	-0.84	Quarterly	15	1.07
5-7	63	0.07	Report	6	1.14
8-10	27	-0.13	Book Chapter	4	0.98
11-14	25	0.51	Book	13	-0.20
15-30	32	0.34			
31-100	53	0.45	<i>Author Gender</i>	<i>Count</i>	<i>Deviation</i>
101-300	53	0.89	Anonymous	119	-0.87
>300	59	1.24	Female	56	-0.29
			Male	374	0.31
<i>Words</i>	<i>Count</i>	<i>Deviation</i>	Male and Female	6	0.77
< 500	180	-0.62			
500-1000	228	-0.06	<i>Media Type</i>	<i>Count</i>	<i>Deviation</i>
1000-2000	86	0.88	Wire	35	-0.67
2000-50,000	48	1.12	Print	369	-0.19
>50,000	13	-0.20	Broadcast	38	0.02
			Web Only	81	0.78
<i>Opinion Level</i>	<i>Count</i>	<i>Deviation</i>			
Editorial	60	-1.31	<i>Specialty</i>	<i>Count</i>	<i>Deviation</i>
Not News Or Editorial	278	0.16	Finance/Business	85	0.84
News	217	0.16	Science/Tech	50	1.18
<i>Politics</i>	<i>Count</i>	<i>Deviation</i>	<i>Misc. Indicator</i>	<i>Count</i>	<i>Deviation</i>
ADA > mean	83	0.10	Cite Insider	92	1.24
ADA < mean	18	0.38	Top Awards	117	0.10
			Top Circulation	64	0.26

Table 1: Average Rating Deviation by Length, Time, Period, Gender, Topic, Etc.




Figure 3: Favorability Rating Vs. Words in Article

and weak, individually increase the favorability toward PAM. There is only one possible exception: books. While books are very long and are often written months before any deadline, books give an unfavorable impression of PAM. But while one (favorable) book devotes 4500 words to related issues, the other twelve books in the sample only discuss PAM in passing, with an average of only about 250 words, or 0.25% of the book, discussing any remotely related topics. Thus the PAM discussion in these books is more like a very short article in a long newspaper.

Looking at the six other variables, we find that males liked PAM more than females, and that web and broadcast articles liked PAM more than print and wire articles. The wire effect, however, is likely due to the fact that almost all the wire articles occurred in the first few days after the Senators' complaint.

To examine the combined effect of all of these information indicators on article favorability, we constructed an (ordinary least squares) multiple linear regression model. The dependent variable is article favorability to PAM, relative to the neutral rating of 4.0. The independent variables are a constant, five other control variables, and the eleven possible information indicators mentioned above, five strong and six weak.¹⁴ As is common when values vary by several orders of magnitude, we took the (base ten) logarithm of (the absolute value of) the three numerical variables: words, days, and period.

¹⁴Note that this regression does not include variables to indicate a wire report, or to indicate whether an ADA score estimate was available for an article. If such variables are added, they are not significant at the (two-tailed) 20% level.

<i>Variable</i>	<i>Coefficient</i>	<i>Std. Err.</i>	<i>% Significance</i>
Log Days	0.64	0.08	$< 10^{-11}$
Insider?	0.87	0.16	$< 10^{-5}$
Finance?	0.61	0.16	$< 10^{-2}$
Science?	0.79	0.21	0.01
Log Words	0.42	0.16	0.42
Editorial?	-0.64	0.25	0.46
News?	0.19	0.13	6.3
Many Readers?	0.46	0.25	3.1
Many Awards?	0.26	0.18	6.7
Log Period	-0.16	0.13	10.7
Anonymous?	0.00	0.19	50
Book?	-1.55	0.47	0.06
Female?	-0.60	0.19	0.07
Web only?	0.43	0.17	0.54
Broadcast?	0.49	0.22	1.5
ADA 10pts	-0.35	0.17	2.2
Constant	-2.81	0.44	$< 10^{-7}$
$R^2 = 0.39$			

Table 2: Multiple Regression of Favorability on Information Indicators

Table 2 shows the parameters of that regression model, giving for each variable the coefficient, the standard error, and the percent threshold for saying the coefficient is significantly different from zero. The strong and weak information indicators, and the other control variables, are shown in order in three different sections of the table. Within each section, variables are ordered by significance.

For the three “Log” variables, the coefficient gives the predicted effect (i.e, the added rating) of a factor of ten increase in that variable. For the binary (or “dummy”) variables (denoted by question marks), the coefficient gives the effect of changing the answer to that binary question from “no” to “yes.” The coefficient of the “ADA 10pts” variable gives the effect of increasing the ADA score by ten points, which is roughly one standard deviation among media source scores.

We had clear expectations about the sign of the relation between all of the information indicators and PAM favorability. We also had a clear expectation about the sign of the ADA score effect. For these variables one-tailed significance percentages are relevant, and Table 2 shows such values. Two-tailed significance percentages are exactly twice as large as the numbers shown; these are the values relevant for the other variables, where we had only weak sign expectations.

For *all* of the strong information indicators, and all of the control variables, the coefficients are in the expected direction, and are significantly different from zero at the 5% level. (That is, all % significance thresholds are less than 5.0.) These coefficients are also substantial, ranging from four to nine tenths of a rating point. Of the six weak information indicators, two are clearly significant at the five percent level, and another two are significant at the seven percent level, and one more at the eleven percent level. The seven percent significant or better coefficients range from two to six tenths of a point, and all are in the expected direction.¹⁵

The model estimates that a 1500 word editorial on July 30, 2003 (i.e., two days later) that mentioned no insiders would give a solidly unfavorable PAM impression (a rating of 2), as would a 50 word non-editorial. But a 60 word news article in an award-winning widely-read science and business publication a month later that mentioned an insider would give a solidly favorable impression (a rating of 6), as would a similar 1000 word article today in a general topic publication. A neutral rating (of 4) is predicted for a 600 word news article in an award-winning widely-read finance publication two days later that does not mention an insider, as does a similar 300 word article a month later in a general-topic publication. (These are all for average-gender identified authors in a daily print publication.)

Discussion

Given the limited number of articles for which we had political estimates, we have probably not fully controlled for political effects. Our estimate of the effect of politics should be of roughly the right magnitude, however, and if so this effect seems too weak to account for

¹⁵Very similar results follow if we just look at each rater’s individual scores.

most of the other strong effects we see.

The large negative book effect can be accounted for in part by noting that many books are really more like a newspaper with a collection of very short articles. In addition there is likely an unmeasured political effect here; many of the twelve books were very political, including the “I Hate George W. Bush Reader,” “50 Reasons Not to Vote for Bush,” “Dude, Where’s My Country,” and “An Exception to the Rulers.”)

The strongest and clearest effect we see is from the passage of time. This is reassuring because this is the variable for which we had the strongest prior expectations of its being an information indicator. In addition, we see that the seven clearly significant information indicators, and both of the marginally significant indicators, all point in the same direction, favoring PAM. This fact seems to lend support to the hypotheses that these supposed information indicators are in fact actually indicators of article information levels regarding PAM, and that more informed articles tended to favor PAM.

The weak finding that publications with longer periods tended to be less favorable is surprising, especially given that this is a positive individual indicator. As the longer period publications are the more academic, and as academia is highly skewed politically (Rothman, Lichter, & Nevitte, 2005), perhaps this reflects an unmeasured political effect.

It is interesting to note that while anonymity was a positive individual indicator, this effects is small or non-existent once we control for lead editorials. Apparently, the main expressions of views that are poorly informed because they are opinionated or anonymous are found in such editorials.

It is also interesting to note that publication prestige, in terms of awards, seems only a weak indicator of being informed about PAM. A stronger indicator seems to be media type; web and broadcast favored PAM more than print, which favored PAM more than books. Apparently, the more recently introduced media types favored PAM more. This adds one more noteworthy result to the conflicting literature on which media types are better informed. The result on gender here can also be added to the controversial literature on which gender tends to be better informed.

Conclusion

When a controversy erupts in the media, and widely differing views are expressed, it is natural to wonder which opinion is the one more favored by those who are most informed about the topic.

There are many standard features which may at least weakly indicate how informed an article is. When there are many media articles on a topic, it is thus possible to statistically study whether these indicators are correlated with each other on a topic, and which side the indicators tend to favor.

This paper describes such an analysis of over five hundred media articles mentioning the Policy Analysis market. Here plausible information indicators are time after the event, article length, whether the article is in a news or an editorial style, whether the author is anonymous, whether the article cites anyone with firsthand knowledge of the event, and the

period, awards, circulation, and topic specialty of the periodical. It turns out that these all individually indicate a more favorable impression of PAM, and when combined in a multiple regression the six clearly significant ones still favor PAM.

This result suggests that while uninformed opinion disliked PAM, informed opinion favored it. Unfortunately, public policy regarding PAM continues to reflect the uninformed opinion, a situation that seems unlikely to change anytime soon. This should give pause to those who think that our democratic processes typically produce informed policy, even when the electorate is highly uninformed. In the PAM case, as in many others (Dixit, 1997; Stiglitz, 1998), policy reflects the uninformed position.

The high level of correlation observed among the information indicators for this topic, with all of the indicators individually pointing in the same direction, is surprising and encouraging. If such information indicators are also highly correlated on other topics, we would have a new and powerful tool both for evaluating both the relative informativeness of different information sources, and for identifying in real time which side in future controversies was better informed about the topic at hand.

References

- Berry, F. C. (1967). A Study of Accuracy in Local News Stories of Three Dailies. *Journalism Quarterly*, 44, 482–90.
- Burriss, L. L. (1985). Accuracy of News Magazines as Perceived by News Sources. *Journalism Quarterly*, 62(4), 824–27.
- Caplan, B. (2001). What Makes People Think Like Economists? Evidence on Economic Cognition from the Survey of Americans and Economists on the Economy. *Journal of Law and Economics*, 44(2), 395–426.
- Clifton, H. (2003). Harsh Criticism Stymies Future of Pentagon's Futures Market. *PR Week*.
- Delli-Carpini, M. X., & Keeter, S. (1997). *What Americans Know about Politics and Why It Matters*. Yale University Press, New Haven.
- Dixit, A. K. (1997). *The Making of Economic Policy: A Transaction-Cost Perspective*. MIT Press, Cambridge, Massachusetts.
- Dunwoody, S., & Scott, B. T. (1982). Scientists as Mass Media Sources. *Journalism Quarterly*, 59, 52–59.
- Furno, C., Bolton, N., & Heim, B. (2002). Evaluating Web Sites and Other Electronic Sources.. <http://camellia.shc.edu/literacy/tablesversion/lessons/WebSitesEval.pdf>.
- Groseclose, T., & Milyo, J. (2004). A Measure of Media Bias.. <http://www.polisci.ucla.edu/faculty/groseclose/MediaBias.pdf>.

- Hanson, G., & Wearden, S. T. (2004). Measuring Newscast Accuracy: Applying a Newspaper Model to Television. *Journalism and Mass Communication Quarterly*, 81(3), 546–558.
- Kraus, N., Malmfors, T., & Slovic, P. (1995). Intuitive Toxicology: Expert and Lay Judgments of Chemical Risks. *Risk Analysis*, 12(2), 215–32.
- Lacy, S., Fico, F., & Simon, T. F. (1991). Fairness and Balance in the Prestige Press. *Journalism Quarterly*, 68(3), 363–70.
- Lott, J. R., & Hassett, K. A. (2004). Is Newspaper Coverage of Economic Events Politically Biased?.. <http://ssrn.com/abstract=588453>.
- Maier, S. R. (2002). Getting It Right? Not in 59 Percent of Stories. *Newspaper Research Journal*, 23, 10–24.
- Moore, B., & Singletary, M. (1985). Scientific sources' perceptions of network news accuracy. *Journalism Quarterly*, 62(4), 816–823.
- Pask, J. M., Kramer, R. J. K., & Mandernack, S. B. (1993). The Savvy Student's Guide to Library Research..
- Polk, C., Hanson, R., Ledyard, J., & Ishikida, T. (2003). The policy analysis market: an electronic commerce application of a combinatorial information market. In *ACM Conference on Electronic Commerce*, pp. 272–273.
- Rothman, S., Lichter, S. R., & Nevitte, N. (2005). Politics and Professional Advancement Among College Faculty. *BE Press The Forum*, 3(1). <http://www.bepress.com/forum/vol3/iss1/art2>.
- Scanlon, J. T. (1972). A New Approach to the Study of Newspaper Accuracy. *Journalism Quarterly*, 49, 587–94.
- Singer, E. (1990). A Question of Accuracy: How Journalists and Scientists Report Research on Hazards. *Journal of Communication*, 40(4), 102–116.
- Singletary, M. (1980). Accuracy in News Reporting: A Review of the Research. Tech. rep. 25, ANPA News Research Report.
- Smardz, Z. (2005). Just Give It a Shot, Girls. *Washington Post*. <http://www.washingtonpost.com/ac2/wp-dyn/A2538-2005Mar26>.
- Stiglitz, J. (1998). The Private Uses of Public Interests: Incentives and Institutions. *Journal of Economic Perspectives*, 12(2), 3–22.
- Tucker, J., & Meirowitz, A. (2004). Learning from Terrorist Markets. *Perspectives on Politics*, 2(2), 331–6.