

Is Any Publicity Good Publicity? A Note on the Impact of Book Reviews

Alan T. Sorensen* and Scott J. Rasmussen†

April 2004

Abstract

This paper uses detailed weekly data on sales of hardcover fiction books to evaluate the impact of New York Times book reviews on sales. In order to weigh the relative propensity of reviews to inform and to persuade, the analysis utilizes a measure of review opinion obtained through a systematic reading of each review. The estimates indicate that in the case of book reviews, any publicity is good publicity: even negative reviews lead to increases in sales. We interpret this finding as evidence that book reviews serve largely to inform consumers about books' content and characteristics (including the books' existence). However, positive reviews have a larger impact on sales than negative reviews, suggesting that reviews also have a persuasive effect.

1 Introduction

When the quality of a product is unknown prior to consumption, product reviews by experts or other consumers can play an important role in shaping demand. The prevalence of published product reviews (which can be easily found for anything from automobiles to preschools) may be taken as presumptive evidence that they affect demand, but basic questions about the type of information they transmit still remain. Do reviews persuade consumers to change their opinions about a product, or do they simply provide information about the product's characteristics (including the fact that it exists)?

*Stanford University and NBER; asorensen@stanford.edu.

†Stanford University; scott.rasmussen@stanfordalumni.org We wish to thank Jim King and Nielsen BookScan for providing the data, and Nanci Moore for assistance in reading the reviews.

This paper addresses these questions in the specific context of book sales. Using weekly data from a private market research firm, we analyze the impact of New York Times book reviews on the sales of 175 hardcover fiction titles published between 2001 and early 2003. The reviews themselves are systematically classified as positive or negative in order to distinguish the effects of good and bad publicity. The basic idea is that if book reviews are primarily persuasive—i.e., their effect is to influence consumers’ evaluations of a book’s quality—then the opinions expressed in the reviews should matter. By contrast, if the reviews are primarily informative—i.e., their effect is to inform readers of a book’s existence and/or characteristics—then any review, good or bad, would tend to increase sales.

Inferring the impact of book reviews by comparing reviewed books to non-reviewed books is obviously problematic, since the New York Times’ decision to review a book may depend on its potential popularity. In order to avoid endogeneity bias, our analysis focuses solely on reviewed books, using time series variation in sales (for a given book) to identify the effect of published reviews. Essentially, we measure the spike in sales in the week immediately following the book review, and ask whether the spike is larger for positive reviews than for negative reviews.

This paper extends a broad literature examining markets in which consumers are uncertain about product quality. Eliashberg and Shugan (1997), Basuroy *et al* (2003), Shaffer and Zettelmeyer (2002), and Reinstein and Snyder (2000) address product reviews in particular; others have analyzed the role of word-of-mouth communication,¹ reputation,² and prices as signals of product quality.³ The questions that motivate this study are also closely related to those that motivate many studies of advertising. For example, Ackerberg (2001) analyzes supermarket sales of yogurt in

¹See Vettas (1997) and Vives (1997).

²See Shapiro (1982), Allen (1984), and Png and Reitman (1995).

³See Nelson (1970), Wolinsky (1983), and Milgrom and Roberts (1986).

combination with data on advertising exposure to quantify the informative vs. persuasive roles of yogurt advertisements. Ahluwalia *et al* (2000) explore the impact of negative publicity on consumer attitudes toward products or brands; some of their idea extend to the present discussion of unfavorable book reviews.

Section 2 presents a basic theoretical framework for understanding how reviews and review quality can influence sales. In Sections 3 and 4 we describe the data, empirical models, and results. The main findings are that book reviews have a positive and statistically significant effect on sales, and that this effect is significantly larger for positive reviews than for negative reviews. We estimate that a positive review leads to a 62.9 percent increase in sales in the week following the review (relative to what sales would have been in the absence of any review), versus a 34.4 percent increase for a negative review. We interpret the difference as evidence that reviews have a persuasive effect on consumers. However, the positive impact of negative reviews suggests (subject to some caveats) the purely informative effect is at least as large as the persuasive effect.

2 A simple theoretical framework

The distinction between the informative and persuasive roles of book reviews is best illustrated in a probabilistic model of book demand. Suppose that a consumer will buy a book only if her valuation of the book (e.g., her interest level) exceeds the price, but that she cannot buy a book unless she is first informed about it. Some consumers (including some whose valuations would exceed the price) will not buy, either because they have too little information about the book's content or because they are completely unaware of the book's existence. Described this way, expected sales can be written as Nq , where N is the number of informed consumers, and q is the probability that

a consumer is interested enough to buy the book given that she is informed.

The impact of a book review depends on how it changes both N and q : after a review, expected sales are $(N + \Delta N)(q + \Delta q)$, with ΔN being the “informative” effect of the review and Δq being the “persuasive” effect. Consider first the informative effect. Many consumers who read the review may already have been informed about the book, so the change in N comes only from previously uninformed consumers who happen to read the review. Note that for well-established authors, many (if not most) potential buyers will be aware of the author’s new book well before any review is published. For relatively unknown authors, however, there may be many consumers who were completely unaware of a book prior to reading its review—i.e., ΔN may be quite large relative to N for obscure authors. This suggests that if book reviews have a significant informative effect, then the impact of a review on sales should be most dramatic for books by relatively unknown authors.

The persuasive effect of a book review is reflected in Δq , the change in the probability that an already-informed consumer will want to buy the book. Note that most of the informed consumers will be infra-marginal, in the sense that even a sparkling review couldn’t convince them to buy the book (or, for others, even the most trenchant critique couldn’t convince them not to buy). For those that are near the margin, however, positive reviews may sway them in favor of purchasing the book, whereas negative reviews may deter them from purchasing. Therefore, to the extent that reviews have a persuasive effect on consumers, positive reviews should have a significantly larger positive impact on sales than negative reviews.

Finally, note that in the absence of any informative effect, negative reviews might be expected to reduce sales. However, in the presence of both informative and persuasive effects, the net impact

of a negative review on sales is ambiguous. If the net effect is always positive, one could rightly say that when it comes to book reviews, “all publicity is good publicity.” An important caveat in our particular case is that reviewed books may be a selected sample. The mere fact that a book is reviewed may signal that it is an outstanding book relative to the universe of books available; negative reviews may simply indicate inferiority relative to the best *reviewed* books. In that case, all reviews, positive or negative, would be expected to increase sales.

3 Background and Data

The dataset analyzed here consists of weekly national sales for 175 hardcover fiction titles that were released from 2001 to 2003 and reviewed by the New York Times.⁴ DMA-level data were also obtained for a subsample of 33 books published in early 2003. The sales data were provided by Nielsen BookScan, a market research firm that tracks book sales using scanner data from an almost-comprehensive panel of retail booksellers.⁵ Additional information about individual titles (such as publication date, subject, and author information) was obtained from a variety of sources, including Amazon.com and the volunteer website Overbooked.org.

Unlike movie critics, book reviewers do not use stars or thumbs-up/thumbs-down systems to summarize their opinions, so we had to do it for them. In order to avoid subjective biases, we attempted to do this in a systematic way. Typical reviews consist primarily of non-opinionated prose describing the book’s characters and plot, so we flagged the sentences likely to be opinionated by using a textual search algorithm (looking for keywords such as the author’s name or the

⁴Only full-length reviews are included in the sample; paragraph-length reviews and articles mentioning multiple books are omitted.

⁵BookScan collects data through cooperative arrangements with virtually all the major bookstore chains, most major discount stores (like Costco), and most of the major online retailers (like Amazon.com). They claim to track at least 80 percent of total retail sales.

word “writing”), randomized them, and then judged them individually as either positive, negative, or neutral.⁶ Each potentially opinionated sentence is therefore evaluated on its own merit. The relative opinion score for each book was then calculated as the ratio of positive sentences to opinionated sentences; for example, a book with P positive sentences, N negative sentences, and Z neutral sentences received a score of $P/(P + N)$. In the econometric analysis, reviews are categorized as positive if this ratio exceeds 67%, negative if it is below 33%, and neutral otherwise. Summary statistics for the sentence and opinion data are reported in Table 1.

The primary motivation for this approach, which essentially involves removing opinionated sentences from the larger context of the review, was to avoid confounding subjective opinions about a book’s content with objective evaluations of the reviewer’s critique. For example, a reader enthused about the history of the American West may believe that a review of a Western novel is positive, even though the reviewer criticizes the novel as dull or otherwise poorly written. We expect our measure to differ somewhat from (but still be positively correlated with) scores based on an individual’s subjective reading of the reviews;⁷ our intent is to approximate the average subjective reading in the population (e.g., averaging over readers who do and don’t like Westerns). A potential drawback to our method is that it doesn’t give extra weight to extreme expressions of opinion: e.g., a review with five mildly positive sentences and one sentence saying “this is the worst novel I’ve ever had the misfortune to read” will still be characterized as a positive review.

Based on our reading, we are not too concerned about this problem: reviewers who write for the

⁶The details of the Perl script we used to find the opinionated sentences are available on request. The algorithm prioritized Type II rather than Type I error minimization in order to limit the number of opinionated sentences excluded from the analysis: 42.3 percent of flagged sentences were opinionated, whereas among a sample of 200 unflagged sentences only 18 (9 percent) were opinionated.

⁷A separate research assistant read a subset of 30 reviews and scored them on a five-point scale from negative to positive. The correlation between her scores and our measure was 0.44 (statistically significant at the 2% level).

New York Times are given to relatively nuanced prose, and blunt expressions of extreme opinion are rare.

4 Empirical Results

Because the decision to review a book may itself depend on the book’s quality, inferring the impact of reviews by comparing reviewed books to non-reviewed books is problematic. Instead, the analysis here focuses solely on reviewed books, identifying the impact of reviews from the time series variation in sales for a given book. We restrict our attention to the nine-week window surrounding each book’s review week.⁸ As indicated in Table 1, most books are reviewed shortly after their release. In some cases, reviews are published within one week of a book’s release; we were forced to omit such books from our sample, since we would have no pre-review sales trajectory to use as a benchmark for measuring the impact of the review.

Suppose that sales of book i in week t are given by

$$s_{i,t} = s_{i,t-1} \cdot \exp\{x'_{i,t}\beta\} \cdot \epsilon_{i,t} \quad (1)$$

so that

$$\ln\left(\frac{s_{i,t}}{s_{i,t-1}}\right) = x'_{i,t}\beta + u_{i,t} \quad (2)$$

We take equation 2 as the model to be estimated. The dependent variable is the log of the sales ratio: today’s sales as a fraction of yesterday’s sales. Note that one advantage of this specification is that it puts all books on the same scale, even though the sales levels vary dramatically across books in the

⁸The window includes the week of the review, plus up to four weeks pre- and post-review.

sample. Another motivation for this specification relates to its predictive power: since measuring the impact of book reviews requires knowing what sales would have looked like in the absence of a review, we want a model that does a good job predicting the path of sales. Allowing sales to depend on lagged sales generates predicted sales paths that are smoother and more accurate than what would be obtained from a model containing only contemporaneous explanatory variables. Also, note that since we use time series variation to identify the impact of reviews, in order for endogeneity bias to be a problem it would have to be the case that reviews are systematically timed to correspond with large unobserved demand shocks—which is a relatively implausible scenario.

Our vector of covariates, $x_{i,t}$, contains variables related to book reviews (e.g., an indicator for whether a review of book i appeared in week t) as well as week dummies (one for each of the 110 weeks represented in the sample, to control for time trends and seasonal variation in book demand), weeks since the book was released, and indicators for whether the book was announced as a television book club pick. The errors, $u_{i,t}$, are assumed to be independent across books but potentially heteroskedastic across books and potentially dependent over time for a given book.⁹

Table 2 reports the estimated coefficients for five specifications based on equation 2. The first column shows the simplest specification, in which $x_{i,t}$ contains an indicator for whether book i was reviewed in week t .¹⁰ The coefficient is positive and precisely estimated; the point estimate suggests that sales are 45.5 percent higher after a review appears than they would have been without the review. Column two of the table looks at the timing of the sales spike more carefully: according to the estimates, sales increase slightly in the week leading up to the published review, substantially

⁹In the reported results, we simply present standard errors that are robust to the potential heteroskedasticity and within-group dependence. Testing directly for serial correlation in the errors is not a trivial exercise in this context; however, when the models are estimated assuming that u follows an AR(1) process, the results are largely unaffected.

¹⁰Specifically, sales in the week where the review indicator equals one represent sales for the Monday-Sunday following the Sunday publication of the review.

more in the week immediately following the review's publication, and then much of the increase in sales disappears after that. The statistically significant increase in sales prior to the appearance of a review could raise doubts about our interpretation of the measured effects as causal. However, this result most likely reflects a quirk in the distribution of the New York Times: the Sunday Book Review is delivered a few days early to mail subscribers.

The third column of the table investigates the differential impact of positive vs. negative reviews. Based on our (continuous) measure of reviews' opinions, we categorize reviews as positive (61), neutral (27), or negative (87), and estimate a separate coefficient for each group. The estimates indicate that positive reviews do have a stronger effect on demand, but that even negative reviews cause sales to increase: a positive review generates a 62.9 percent increase in demand, whereas a negative review leads to a 34.4 percent increase. A standard F -test rejects the hypothesis that the positive and negative review coefficients are equal ($F_{1,174} = 5.32, p = .022$). The differential impact of positive vs. negative reviews implies that reviews are at least partly persuasive: more consumers are convinced to purchase a positively reviewed book than a negatively reviewed one.

One potential alternative explanation for the difference between the coefficients for positive and negative reviews is that positive reviews may be published more conspicuously in the paper. For example, editors may be more inclined to use positive reviews as features, placing them at the front of the Sunday Book Review, so that positive reviews are more likely to be read (and therefore more likely to influence sales) than negative reviews. Also, positive reviews may be given more space, so that they are longer and contain more information. Although we have no data indicating whether the placement of reviews differs systematically depending on the opinions expressed, the

available data reveal no relationship between opinion and review length: the correlation between the number of sentences in a review and the fraction of opinionated sentences that were positive is in fact slightly negative.

While the differential impact of positive vs. negative reviews suggests the presence of a persuasive effect, the measured impact of negative reviews reveals a substantial informative role as well. Even when a reviewer clearly pans a book, the publication of the review leads to an increase in the book's sales. As mentioned in section 2, to the extent that reviews are primarily informative, they should have a larger impact on sales of books by relatively unknown authors. Column 4 of the table reports results from a specification that tests this proposition. The impact of a book review is estimated separately for three categories of authors: those that have published fewer than four books prior to the book in question, those that have published between four and nine, and those that have published ten or more. As expected, the coefficient is largest for the relatively new authors. Taken at face value, these results are consistent with the idea that an important role of reviews is simply to inform consumers (either of a book's existence or of its characteristics).¹¹ However, the relative imprecision of the estimates prevents us from drawing any strong conclusions.¹²

Although the New York Times is read widely throughout the nation, its readership is still somewhat concentrated in New York City. For a small subsample of 33 books, we were able to obtain DMA-specific sales data to test whether the impact of New York Times book reviews is largest in New York City.¹³ The last column of Table 2 reports the estimated effect of book reviews in New

¹¹These results are also consistent with the arguments of Chandy *et al* (2001) that expert sources are more influential in new markets because customers' greater uncertainty about product quality in new markets motivates them to find (and carefully process) information.

¹²An F -test of the hypothesis that the three coefficients are equal fails to reject at the 5 percent level ($F_{2,174} = 2.71$, $p = .0693$). Somewhat surprisingly, specifications in which the effect is separate for debut authors (0 previously published titles) show no difference between debut authors and authors with 1 or 2 previous titles.

¹³Nielsen defines DMA's (designated market areas) for purposes of studying geographic variation in sales. The DMA's generally encompass entire metropolitan areas: e.g., the New York, NY DMA includes Brooklyn and Queens

York City vs. all other DMA's combined. As expected, the effect is much larger in New York, lending some additional credibility to our interpretation of these estimates as causal effects.¹⁴

5 Discussion

How much is a review in the New York Times worth? The point estimates reported in Table 2 allow us to make some back-of-the-envelope calculations, although we obviously regard them as somewhat speculative. Since the empirical models measure the percentage increase in unit sales following a review, the expected increase in revenues depends on both the price and the level of sales prior to the review. In Table 3 we provide estimates of revenue increases for three potential books (at approximately the 25th, 50th, and 75th percentiles of sales levels), assuming a price equal to the sample median of \$24.95. While the numbers aren't large for small-scale books (on the order of \$1,000 of extra revenue for the week after the review's publication), for popular books by established authors the difference between a positive and negative review (or no review at all) can be quite substantial.

The fact that even negative reviews lead to increased sales suggests that indeed all publicity is good publicity. We attribute this finding to an informative effect, whereby reviews serve to inform consumers of a book's existence and content, but we note that the significance of this effect in our data may be peculiar to the market for books. Thousands of new books are published each year in the U.S., so consumers can't possibly be fully informed about the complete set of available books.¹⁵ Consequently, book reviews are very likely to contain information that is new to most

and Manhattan, etc.

¹⁴In unreported specifications we allowed the effect to differ for each of the ten largest DMA's. Although the rank ordering of the coefficients makes sense (the effect in New York was largest, followed by Philadelphia; effects for San Francisco and Detroit were substantially smaller), their magnitudes were difficult to explain in some cases (e.g., the effects in Philadelphia and New York were statistically indistinguishable).

¹⁵The *Bowker Annual* (2001) reports that roughly 110,000 new titles were published in 2000, over 4,200 of which

readers. By contrast, in the motion picture industry only a few hundred movies are released each year. Given the smaller number of products (and the amount of pre-release advertising typically associated with the products), consumers are much more likely to be aware of the basic content of a movie than of a book, so one might expect the relative importance of persuasive effects to be greater for movie reviews. In fact, in their study of movie reviews Reinstein and Snyder (2000) find relatively large “influence” effects (positive reviews have a significantly larger impact on sales than negative or mixed reviews), but find no evidence that sales increase after negative reviews, and Basuroy *et al* report that negative reviews significantly *decrease* box office revenues.

Although our finding of a substantial informative effect may not generalize to markets with higher levels of product awareness, the evidence of a persuasive effect has implications that extend beyond the market for books. Many economists have been reluctant to accept the notion that consumer preferences can be changed or influenced by advertising, instead postulating a number of alternative rationales for observed levels of advertising expenditures. Our results offer direct empirical evidence suggesting consumers’ opinions about a product are malleable, which offers some validation of the most obvious motive for advertising. More broadly, the evident influence of product reviews indicates that learning about product quality is an economically important problem for consumers, and that the efficiency with which a market operates may depend on how easily product information can be transmitted.

were hardcover fiction.

References

- Ackerberg, D. (2001, Summer). Empirically distinguishing informative and prestige effects of advertising. *RAND Journal of Economics* 32(2), 100–118.
- Ahluwalia, R., R. Burnkrant, and H. R. Unnava (2000, May). Consumer response to negative publicity: The moderating role of commitment. *Journal of Marketing Research* 37(2), 203.
- Allen, F. (1984). Reputation and product quality. *Rand Journal of Economics* 15, 311–327.
- Basuroy, S., S. Chatterjee, and S. A. Ravid (2003, October). How critical are critical reviews? the box office effects of film critics, star power, and budgets. *Journal of Marketing* 67(4), 103.
- Bogart, D. (Ed.) (2001). *The Bowker Annual* (46 ed.). R.R. Bowker.
- Chandy, R., G. Tellis, D. MacInnis, and P. Thaivanich (2001, November). What to say when: Advertising appeals in evolving markets. *Journal of Marketing Research* 38(4), 399.
- Eliashberg, J. and S. M. Shugan (1997). Film critics: Influencers or predictors? *Journal of Marketing* 61, 68–78.
- Milgrom, P. and J. Roberts (1986). Prices and advertising signals of product quality. *Journal of Political Economy* 94(796-821).
- Nelson, P. (1970). Information and consumer behavior. *Journal of Political Economy* 78, 311–329.
- Png, I. P. L. and D. Reitman (1995). Why are some products branded and others not? *Journal of Law and Economics* 38, 207–224.
- Reinstein, D. A. and C. M. Snyder (2000, September). The influence of expert reviews on consumer demand for experience goods: A case study of movie critics. Working paper, George Washington University.
- Shaffer, G. and F. Zettelmeyer (2002, Summer). When good news about your rival is good for you: The effect of third-party information on the division of channel profits. *Marketing Science* 21(3), 273.
- Shapiro, C. (1982). Consumer information, product quality, and seller reputation. *Bell Journal of Economics* 13, 20–35.
- Vettas, N. (1997). On the informational role of quantities: Durable goods and consumers' word-of-mouth communication. *International Economic Review* 38, 915–944.
- Vives, X. (1997). Learning from others: A welfare analysis. *Games and Economic Behavior* 20, 177–200.
- Wolinsky, A. (1983). Prices as signals of product quality. *Review of Economic Studies* 50(647-658).

Table 1: Summary statistics for reviews ($n=175$)

	Mean	Std. Dev.	Percentiles		
			.10	.50	.90
Books:					
List price	\$24.63	\$3.12	\$23.00	\$24.95	\$25.95
Average weekly sales	2,685.5	10,831.9	48.9	368.3	3,570.2
Week reviewed (wrt release)	7.7	8.2	3.0	5.0	12.0
Reviews:					
Number of sentences	45.5	22.6	12.0	48.0	74.0
Number opinionated	6.8	4.4	2.0	6.0	13.0
Percent opinionated	16.0%	8.9%	7.1%	14.6%	25.0%
Percent positive	55.3%	30.3%	11.1%	52.6%	100.0%

Table 2: Regression estimates: the impact of reviews on sales

	I	II	III	IV	V
Weeks since release	.004 (.003)	.004 (.003)	.003 (.003)	.004 (.003)	-.014 (.011)
Review	.375 (.042)	–	–	–	–
Review ($t-2$)	–	.064 (.052)	–	–	–
Review ($t-1$)	–	.129 (.057)	–	–	–
Review ($t=0$)	–	.369 (.048)	–	–	–
Review ($t+1$)	–	-.145 (.039)	–	–	–
Review ($t+2$)	–	-.045 (.044)	–	–	–
Positive Review	–	–	.488 (.072)	–	–
Neutral Review	–	–	.383 (.090)	–	–
Negative Review	–	–	.296 (.051)	–	–
Review \times (<3 previous titles)	–	–	–	.449 (.059)	–
Review \times (4-9 previous titles)	–	–	–	.264 (.089)	–
Review \times (>10 previous titles)	–	–	–	.295	– (.069)
Review \times (New York City)	–	–	–	–	.812 (.143)
Review \times (all other markets)	–	–	–	–	.439 (.092)
Number of books	175	175	175	175	33
Number of observations	1,268	1,268	1,268	1,268	3,936
R^2	.244	.267	.249	.249	.083

Robust standard errors in parentheses. Each specification also includes a full set of week fixed effects, and indicators for whether the book was announced as a television book club pick. (A full table of results is available from the authors upon request.) The review dummies equal one in the week immediately following the publication of the book review in the Sunday New York Times.

Table 3: Dollar values of book reviews^a

Weekly sales prior to review ^b	Increase in revenues (good review)	Increase in revenues (bad review)
200	\$1,982	\$1,119
500	\$4,954	\$2,798
2,000	\$19,816	\$11,191

^a Based on the estimated coefficients in column 3 of Table 2, assuming the book's retail price equals the sample median (\$24.95).

^b The three sales levels correspond roughly to the 25th, 50th, and 75th percentiles of weekly sales in the sample.