Journal of Wine Economics, Volume 4, Issue 2, Winter 2009, Pages 125–151

Does Advertising Bias Product Reviews? An Analysis of Wine Ratings*

Jonathan Reuter[a]

Abstract

In markets for experience goods, publications exist to help consumers decide which products to purchase. However, in most cases these publications accept advertising from the very firms whose products they review, raising the possibility that they bias product reviews to favor advertisers. To test for biased product reviews, I exploit the fact that, of the two major U.S. wine publications, only *Wine Spectator* accepts advertising. Although the average *Wine Spectator* ratings earned by advertisers and non-advertisers are similar, I find that advertisers earn just less than one point higher *Wine Spectator* ratings than non-advertisers when I use *Wine Advocate* ratings to adjust for differences in quality. However, I find only weak evidence that the selective retasting of advertisers' wines contributes to the higher ratings. Moreover, conditional on published ratings, *Wine Spectator* is no more likely to bestow awards upon advertisers. I conclude that while advertising may influence ratings on the margin, *Wine Spectator* appears largely to insulate reviewers from the influence of advertisers. (JEL Classification: L15, M37)

I. Introduction

In markets for experience goods, publications exist to inform consumers' decisions about which goods to purchase.[1] These publications introduce consumers to the available products and publish product reviews intended to help consumers rank them. However, with notable exceptions like *Consumer Reports*, these publications receive a substantial portion of their revenue from advertisers. This raises the possibility that publications bias their

* This paper began life as a chapter from my 2002 Ph.D. dissertation at MIT. As such, I would like to thank my primary thesis advisers, Glenn Ellison and Sendhil Mullainathan, for their advice and encouragement. I would also like to thank Peter Davis, Denis Gromb, Jerry Hausman, Robin McKnight, Nancy Rose, Chris Synder, Jeff Wilder, Eric Zitzewitz, two anonymous referees, and participants at the MIT Industrial Organization Workshop for helpful comments on this and previous drafts. Finally, I would like to thank Karl Storchmann for soliciting this paper for review for the *Journal of Wine Economics*. Financial support from an NSF Graduate Research Fellowship is gratefully acknowledged. Any errors are my own.

[a] Department of Finance, Carroll School of Management, Boston College, 140 Commonwealth Avenue, Chestnut Hill, MA 02467, email: reuterj@bc.edu.

[1] The utility generated by many goods can only be learned through use. Nelson (1970) terms these experience goods and notes that experience goods are likely to give rise to third–party product reviews.

contents to favor advertisers. Specifically, publications that accept advertising might review disproportionately more products from firms that advertise *(biased product coverage),* or they might review products from advertisers more favorably than comparable products from non-advertisers *(biased product reviews)*. Both forms of bias are likely to benefit advertisers through increased demand.[2] However, whether it is optimal for a publication to bias its content depends on how the introduction of bias would impact the number of consumers who rely upon that publication.

Consider the profit maximization problem for a publication that derives revenue from both advertisers and subscribers.[3] Advertising revenues are increasing in the number of subscribers and the extent to which the publication biases its content to favor advertisers (for a fixed number of subscribers).[4] Subscription revenues are increasing in the number of subscribers and the price of a subscription. Finally, the number of subscribers is decreasing in the price of a subscription and increasing in the publication's reputation for being unbiased. This simple model yields three insights. First, the more costly that it is for consumers to rely upon a biased publication, the less likely the publication is to permit bias. In this case, the increase in advertising revenues associated with bias is likely to be outweighed by the decrease in advertising and subscription revenues associated with fewer subscribers. Second, because demand is increasing in a publication's reputation for being unbiased (at a fixed subscription rate), it is optimal for publications to take actions that prevent—or appear to prevent—advertiser status from influencing content. Third, consumers may be willing to accept bias if the higher advertising revenues that arise from bias are used to lower subscription rates. This last point provides two ways to interpret the continued existence of publications that both accept advertising and review advertisers' products. Either these publications are unbiased or consumers would rather pay the lower subscription rates that advertising bias makes possible than fund the publications entirely through higher subscription rates. It is, therefore, an empirical question whether any particular publication biases its content.

In this paper, I use data collected from two major U.S. wine publications to test for biased product reviews.[5] There are several reasons why wine publications are a nice environment in which to test for potential pro-advertiser biases. First, because each vintage of each wine is a

[2] Reinstein and Snyder (2005), Reuter and Zitzewitz (2006), and Del Guercio and Tkac (2008) provide evidence that expert opinion influences demand for experience goods. However, none of these papers studies how the influence of biased opinion differs from that of unbiased opinion.

[3] This analysis is similar in spirit to Darby and Karni (1973).

[4] Setting aside whether advertising indirectly influences demand through biased content, there are numerous ways in which advertising might directly influence demand. For example, in his seminal article on the information content of advertisements, Nelson (1974) argues that advertisements for experience goods increase demand by signaling product quality. More recently, Ackerberg (2001) presents evidence that advertising for a particular experience good (a new brand of yogurt) informed consumers about the existence of the good rather than increased brand appeal.

[5] In a related study, Reuter and Zitzewitz (2006) study the impact of advertising on mutual fund recommendations. They find that advertisers are more likely to receive positive mentions and less likely to receive negative mentions, but only in the subset of publications most dependent on advertising by mutual fund families. More generally, DellaVigna and Gentzkow (2009) survey the growing empirical literature on persuasion.

different good, the typical issue of a wine publication contains several hundred reviews, which translates into a correspondingly large number of observations when conducting statistical tests. Second, each of the major U.S. wine publications rates wine quality on a scale that ranges from 50 to 100. The fact that each of these publications distills quality to a single number allows for easy comparisons of wine ratings across publications. Third, none of the major U.S. wine publications specializes in wines from a particular region or set of regions, thereby ruling out any obvious differences in taste. Fourth, while consumers presumably have a preference for unbiased wine reviews, biased wine reviews impose low enough costs on consumers that biased product reviews of wine are more likely than, for example, biased product reviews of consumer durables. Fifth, of the two most influential U.S. wine publications, only one accepts advertising. *Wine Spectator* is funded by a mixture of advertising revenue and subscriptions, while *Wine Advocate,* "the independent consumer's bimonthly guide to fine wine," is funded entirely by subscriptions. The fact that *Wine Advocate* uses the lack of advertising as a selling point suggests that some consumers expect advertising to bias wine reviews in other publications.[6]

The claim that a publication biases its product reviews is the claim that advertisers receive more favorable reviews on their products than they would have received if they were not advertisers. Because I do not observe the counter-factual reviews, I use reviews from *Wine Advocate* to proxy for the true quality of the wines reviewed by *Wine Spectator*. Controlling for (censored) *Wine Advocate* ratings, I find a positive (partial) correlation between *Wine Spectator* ratings and measures of lagged advertising intensity. The implication is that *Wine Spectator* ratings of advertisers' wines are approximately one point higher than their ratings of comparable wines from non-advertisers. (The fact that *Wine Advocate* essentially only publishes ratings of 85 and above complicates the tests for biased wine ratings. In Section 3, I describe how to overcome these complications.) To the extent that consumer demand rises with the *Wine Spectator* rating, (slightly) higher ratings should translate into (slightly) higher prices.

Although my finding that advertisers receive higher *Wine Spectator* ratings is consistent with biased wine ratings, it is also consistent with the two publications evaluating wine using different standards, perhaps because they cater to different consumer tastes. Specifically, if *Wine Spectator* systematically favors some types of wines, *Wine Advocate* systematically favors other types of wines, *and* expected differences in *Wine Spectator* and *Wine Advocate* ratings drive decisions about how much to advertise in *Wine Spectator*, the estimated (partial) correlation between *Wine Spectator* ratings and advertising intensity could be positive even in the absence of biased wine ratings. In other words, the higher *Wine Spectator* ratings of advertisers could reflect the fact that wineries who make "*Wine Spectator* style" wines advertise in *Wine Spectator* because their wines are the most likely to appeal to

[6] For example, according to Ashenfelter (1989), "Newspaper writers often accept direct and indirect payoffs for touting wines and many knowledgeable consumers are aware of this practice. (After a detailed expose by media reporter David Shaw in the *Los Angeles Times* in the summer of 1987, the *Times* fired its on wine writer!) The result is that there is a market for independent information. Robert Parker, lawyer turned wine writer, has captured that market and the absolute faith of his readers by refusing to accept payoffs." Robert Parker is the publisher of the *Wine Advocate*. During my sample period, he is also the main reviewer.

Wine Spectator readers.[7] Fortunately, there are several reasons to doubt that expected differences between *Wine Spectator* and *Wine Advocate* ratings drive advertising in *Wine Spectator*. First, wineries that advertise tend to be the large-production wineries, and while *Wine Spectator* may have a preference for reviewing large-production wines, it is not clear that large wineries should be more likely to produce wines whose qualities appeal to *Wine Spectator's* reviewers. Second, both publications maintain that their ratings have objective merit. According to a statement on the cover of each issue of *Wine Advocate*, "While some have suggested that scoring is not well suited to a beverage that has been romantically extolled for centuries, wine is no different from any consumer product. There are specific standards of quality that full-time wine professionals recognize, and there are benchmark wines against which all others can be judged." Table 1 provides a key for interpreting *Wine Spectator* and *Wine Advocate* ratings. Nevertheless, because identification is crucial, I use data on which advertisements cite *Wine Spectator* ratings to test whether advertising by "*Wine Spectator* style" wineries introduces a positive bias into estimates of biased wine ratings. Interestingly, I find a negative (partial) correlation between the fraction of a winery's advertisements that cite prior *Wine Spectator* ratings (or awards) and the ratings that *Wine Spectator* bestows upon its wines.

Table 1
Interpreting *Wine Spectator* and *Wine Advocate* Ratings

	Wine Spectator's 100-Point Scale
95–100	Classic: a great wine
90–94	Outstanding: a wine of superior character and style
85–89	Very good: a wine with special qualities
80–84	Good: a solid, well-made wine
70–79	Average: a drinkable wine that may have minor flaws
60–69	Below average: drinkable but not recommended
50–59	Poor, undrinkable: not recommended
	Wine Advocate's 50–100 Point Scale
96–100	An extraordinary wine of profound and complex character displaying all the attributes expected of a classic wine of its variety
90–95	An outstanding wine of exceptional complexity and character
80–89	A barely above average to very good wine displaying various degrees of finesse and flavor as well as character with no noticeable flaws
70–79	An average wine with little distinction except that it is soundly made
60–69	A below average wine containing noticeable deficiencies
50–59	A wine deemed to be unacceptable

The description of *Wine Spectator* ratings is taken from the November 15, 2000 issue. The description of *Wine Advocate* ratings taken from the introduction to issue 132 (dated December 23, 2000). Both scales are identical to those published in other issues from which data are drawn.

[7] Of course, this problem of identification is not limited to tests for biased wine ratings. If publications cater to consumers with particular tastes, and their product reviews reflect those tastes, then even when ratings are unbiased, firms may choose to advertise predominately in those publications that rate their products most highly.

Above, I stated that the need for credibility with consumers might prompt publications to take steps to separate advertisers from reviewers. In the case of *Wine Spectator*, this is accomplished through the use of blind tastings, in which "tasters are told only the general type of wine (varietal or region) and the vintage." According to a statement that appears in each issue of *Wine Spectator,* "Notes and ratings are entered directly into our database prior to the removal of the bags. Additional comments may be added to a note after the identity of the wine is revealed, but the score is never changed." Blind tastings speak to the value *Wine Spectator* places on maintaining a reputation for unbiased reviews, and would appear to rule out biased wine ratings. However, biased wine ratings could still arise through selective retasting. According to *Wine Spectator*, "We retaste all wines that score 70 points or less. We retaste many other wines to confirm impressions. Outstanding scores are routinely confirmed by another editor." Therefore, if *Wine Spectator* were more likely to retaste wines from wineries that advertise—perhaps because it is easier to obtain additional bottles from advertisers—and if published ratings reflect the quality of the best bottle tasted, these facts will collectively bias wine ratings in favor of advertisers. Because reviews indicate when published ratings are based on multiple tastings of the same wine, I am able to test whether biased wine ratings arise through selective retasting. While I find evidence that *Wine Spectator* is more likely to retaste an advertiser's wine, within the full sample of U.S. wines reviewed by *Wine Spectator*, retastings are associated with lower *Wine Spectator* ratings for advertisers and non-advertisers alike. Moreover, when I control for retastings within this sample, I continue to find that advertisers earn higher ratings than non-advertisers. However, I also find that *Wine Spectator* is more likely to retaste wines rated by *Wine Advocate*. Within this smaller sample of wines, I find that retastings differentially increase the ratings of advertisers. However, the number of retasted wines is much too small to explain a one-point difference in ratings.

Finally, I use the fact that *Wine Spectator* bestows awards upon a small fraction of reviewed wines to test for biased awards. The four awards are *Best Buy*, *Cellar Selection*, *Highly Recommended*, and *Spectator Selection*, which is the publication's highest award. Although the awards are targeted at different types of wines, relative to ratings, *Wine Spectator* should have considerably more discretion in deciding which wines receive awards. Consequently, if *Wine Spectator* biases its content to favor advertisers, one might expect more bias along the awards dimension than along the ratings dimension. As expected, I find that the probability of receiving an award is increasing in *Wine Spectator's* published rating. Therefore, everything else equal, biased ratings should increase the probability of receiving an award. However, conditional on price, production level, and rating, I find that advertisers are either no more likely or less likely to receive awards than non-advertisers.

Differences in the subscription rates, circulations, and contents of the two publications prevent me from testing for biased wine coverage, because these differences suggest that it is optimal for *Wine Spectator* to review fewer small-production wines than does

Wine Advocate.[8] Since the wineries that advertise are more likely to sell large-production wines, it then becomes impossible to identify whether wines are reviewed in *Wine Spectator* because they are widely-available and, hence, of interest to its readers or because they come from wineries that advertise.[9] However, the fact that only 9.8% of the wines reviewed by *Wine Spectator* come from wineries that advertise, certainly suggests that *Wine Spectator* is willing to review wines from wineries that do not advertise.

Overall, the tests for biased ratings and biased awards produce little consistent evidence that *Wine Spectator* favors advertisers. At worst, the tests for biased ratings suggest that *Wine Spectator* rates wines from advertisers almost one point higher than wines from non-advertisers. However, selective retastings can explain at most half of this bias and only within the set of U.S. wines rated by both *Wine Spectator* and *Wine Advocate*. Given *Wine Spectator's* claim that it rates wines blind, the remaining difference in ratings may simply reflect consistent differences in how the two publications rate quality. The fact that tests for biased awards provide no evidence of bias suggests that there is little bias overall. Therefore, despite the fact that *Wine Spectator* is dependent on advertising revenue, the long-run value of producing credible reviews appears to limit bias.

The remainder of the paper is organized as follows. Section 2 describes the data used to test for biased ratings and biased awards and summarizes how wines produced by advertisers differ from those produced by non-advertisers. Section 3 formulates tests for biased ratings and discusses a number of econometric complications. Section 4 presents the empirical results of each test, and Section 5 concludes.

II. Overview of the Data

In this paper, I test for biased wine reviews using a unique database that combines wine review and advertising data from three sources. First, wine reviews published in *Wine Spectator* during 1999 and 2000 were obtained from *Wine Spectator Online*. From these reviews, I extracted the name and vintage of each wine, the name of the winery that produced it, the grape varietal(s) from which it was produced, the region(s) and country in which the grapes were grown, the production level (measured in cases), the U.S. retail price, the *Wine Spectator* rating, the *Wine Spectator* award (if any), and whether the published rating was

[8] During my sample period, *Wine Advocate* charged $55 per year for six issues containing little more than ratings and tasting notes, whereas *Wine Spectator* charged $45 per year for eighteen issues containing wine reviews, restaurant reviews, travel guides, recipes, and glossy color photographs. Moreover, the audited circulation of *Wine Spectator* was close to 325,000 (with an estimated 1.5 million readers) while the audited circulation of *Wine Advocate* was closer to 40,000. According to *Wine Spectator*, "Wines are chosen for tasting among those sent to our offices for review and those purchased at retail. Because we for the most part serve a U.S. audience, we prefer to review wines that are widely-available there and merit wide interest for that readership."

[9] More generally, to claim that a publication biases its product coverage is to claim that *among the set of products that interest its readership* the publication reviews a disproportionate number of products from advertisers. Therefore, testing for biased product coverage requires that one be able to identify which products not reviewed by a publication are of interest to its readers.

the result of multiple tastings. Using wine name, winery name, and vintage variables, I also attempted to determine the *Wine Spectator* rating awarded to the prior vintage of each wine reviewed in 2000. Wines tasted from barrel and wines tasted as part of special tastings (such as a vertical tasting of the wines from a particular winery) were dropped from the sample, reducing the total number of reviews in 2000 from 11,371 to 10,935.[10]

Second, from each of the *Wine Spectator* issues published during 1999 and 2000, I collected data on which wineries advertised in *Wine Spectator*, the page size of each winery's advertisement, and whether the advertisement cited a *Wine Spectator* rating or award.[11] The advertising data from 1999 were used to calculate the number of pages of advertising placed by each winery through the end of 1999, as well as the number of pages of advertising that cited a *Wine Spectator* rating or award. For each wine reviewed in 2000, I classified the wine as coming from an advertiser if the winery advertised at least once in the prior 12 months.

Finally, for each of the 2,753 U.S. wines reviewed by *Wine Spectator* in 2000, I sought a *Wine Advocate* rating. Using issues 116 (dated April 30, 1998) through 135 (June 23, 2001) of *Wine Advocate*, I located ratings for 713 of the 2,753 U.S. wines reviewed by *Wine Spectator*. The other 2,040 wines either were not tasted by *Wine Advocate* or, because *Wine Advocate* rarely publishes rates less than 85 points, were tasted and rated less than 85 points. I did not gather data on U.S. wines reviewed by *Wine Advocate* but not reviewed by *Wine Spectator*.

Table 2 summarizes the two samples used throughout this paper. Panel A focuses on the subset of U.S. wines reviewed by *Wine Spectator* in 2000; this is the sample of wines used to test for biased wine ratings. Panel B focuses on the full set of wines reviewed by *Wine Spectator* in 2000; this larger sample of wines is used to test for biased awards. "WS rating" and "WA rating" correspond to published *Wine Spectator* and *Wine Advocate* ratings, excluding barrel samples and wines reviewed as part of special tastings. "WS prior rating" is the *Wine Spectator* rating for the prior vintage of a particular wine. The fact that the number of "WS prior ratings" is less than half the number of "WS ratings" reflects the facts that *Wine Spectator* does not review the same set of wines every year and that *Wine Spectator* does not always record wine names consistently from one vintage to the next. "WA tasted" equals one if *Wine Advocate* published a review (though not necessarily a rating) for the wine and "WA recommends" equals one if *Wine Advocate* rated the wine 85 points or above. "Price" is the retail price as reported in *Wine Spectator* at the time of the review. "Production" is the production level of the wine, measured in cases. It is available for approximately 97.5 percent of U.S. wines and for approximately 75 percent of all wines. "Ad Pages" measures the number of pages of advertising taken out by each winery in the 18 issues published over the prior 12 months, while "Ad Pages WS" measures the number of these pages that cite a *Wine Spectator*

[10] Wines reviewed from barrel were dropped because they receive a ratings range like 85–89 rather than an integer rating. Wines reviewed as part of a special tasting were dropped because they almost always represent a second review of a previously–rated wine. Moreover, neither type of review is eligible to receive a *Wine Spectator* award.

[11] Because there was no systematic way to do so, I did not gather information on which wineries advertised on *Wine Spectator Online*. I discuss limitations of the advertising data in Section 4.

Table 2
Summary Statistics for Wines Reviewed by *Wine Spectator* in 2000

Variable	*Obs.*	*Average*	*Std. Dev.*	*Minimum*	*Maximum*
		Panel A: U.S. Wines			
WS rating	2,753	85.894	4.028	68	99
WS prior rating	1,295	87.520	3.675	71	99
WA rating	705	88.556	3.291	64	100
WA tasted	2,753	0.259	0.438	0	1
WA recommends	2,753	0.247	0.431	0	1
Price	2,743	28.350	21.939	6	395
Production (cases)	2,683	8,522.3	46,956.2	24	1,000,000
Ad Pages	2,753	0.629	1.759	0	13
Ad Pages *WS*	2,753	0.117	0.686	0	8
Ad Dummy	2,753	0.187	0.390	0	1
Ad Dummy *WS*	2,753	0.045	0.207	0	1
Ad Ratio	2,753	0.033	0.166	0	1
Retasted Dummy	2,753	0.044	0.205	0	1
		Panel B: All Wines			
WS rating	10,935	85.388	4.718	55	100
Price	10,331	31.848	35.621	5	1,000
Production (cases)	8,156	7,094.0	32,876.7	5	1,000,000
Ad Dummy	11,213	0.098	0.297	0	1
Ad Pages	11,213	0.349	1.495	0	14

Panel A is restricted to U.S. wines reviewed by *Wine Spectator* in 2000; Panel B includes wines produced in other countries. Wines reviewed from barrel or as part of special tastings are excluded. The number of *WA* ratings is less than given in Table 3 because eight wines are simply rated "?".

rating or award. Measured in this way, less than one in five pages of advertising cites *Wine Spectator*. "Ad Dummy" equals one if the winery advertised in *Wine Spectator* in the prior 18 issues (*i.e.*, if "Ad Pages" is positive), and "Ad Dummy WS" equals one if any of a winery's advertisements cited *Wine Spectator*. "Ad Ratio" is the ratio of a winery's advertisements that cite *Wine Spectator* to those that do not (and equals zero when "Ad Pages" equals zero). Finally, the "Retasted Dummy" equals one if the published *Wine Spectator* rating was based on multiple tastings of the wine.[12] Selective retasting provides a mechanism through which biased wine ratings might arise, although the extent of possible bias is limited by the fact that only 4.4 percent of wines are reported as having been retasted.

The distributions of *Wine Spectator* and *Wine Advocate* ratings for the sample of U.S. wines reviewed by *Wine Spectator* in 2000 are compared in Figure 1. Panel (a) contains a histogram of *Wine Spectator* ratings for those wines for which *Wine Advocate* ratings are

[12] For example, the retasted dummy equals one for the 1997 Beaulieu Vineyard Cabernet Sauvignon Rutherford Clone 6 because *Wine Spectator* included the following phrase in its description of the wine: "This was the best of four bottles tasted, with the other three exhibiting signs of cork taint." The retasted dummy variable is unrelated to whether *Wine Spectator* has published other reviews of the same vintage of the same wine.

Figure 1
Wine Spectator and _Wine Advocate_ Ratings

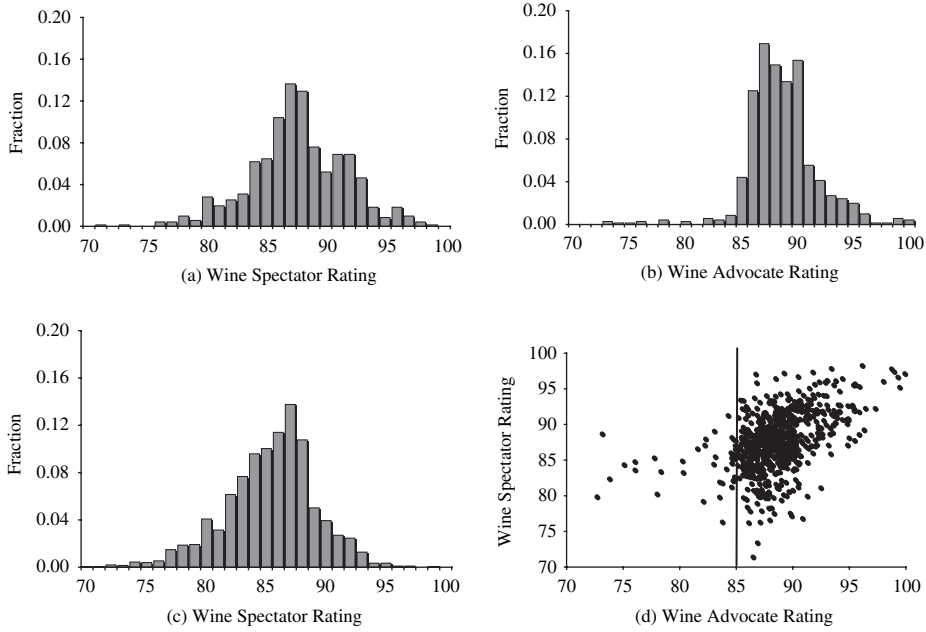

Panel (a) contains the distribution of *Wine Spectator* ratings for which *Wine Advocate* ratings are also observed; panel (b) contains the distribution of *Wine Advocate* ratings; panel (c) contains the distribution of *Wine Spectator* ratings for which *Wine Advocate* ratings are not observed; and panel (d) plots *Wine Spectator* ratings versus *Wine Advocate* ratings. In each case, the sample is restricted to wines reviewed by *Wine Spectator* in 2000.

available and panel (b) contains a histogram of the *Wine Advocate* ratings, in which it is apparent that *Wine Advocate* rarely publishes ratings less than 85 points. Panel (c) contains a histogram of *Wine Spectator* ratings for the set of U.S. wines not reviewed by *Wine Advocate*. Comparing panels (a) and (c) reveals that the sample of wines for which *Wine Advocate* ratings are available is a selected sample of all U.S. wines reviewed by *Wine Spectator.* Panel (d) contains a scatter plot of *Wine Spectator* ratings and matched *Wine Advocate* ratings.[13] The lack of *Wine Advocate* below 85 points is apparent, as is the relatively low correlation (0.493) between the ratings published in the two publications. To the extent that the publications measure quality using the same standards, this low correlation suggests that one or both ratings contain a substantial amount of noise.

Table 3 compares the ratings and characteristics of various subsets of wines. Panel A focuses on U.S. wines and indicates that 18.7 percent of U.S. wines come from wineries that

[13] *Wine Spectator* and *Wine Advocate* ratings are restricted to integers. To capture the number of observations with each pair of ratings in Figure 1(d), I added noise (uniform on the interval between −0.5 and 0.5) to each rating.

Table 3
Characteristics of Wines Reviewed by *Wine Spectator* in 2000

Reviews	*Obs.*	*WS Rating*	*WA Rating*	*Price*	*Production*	*% Advert.*
Panel A: U.S. Wines	2,753	85.89	–	$28.35	8,522	18.7%
WS Only	2,040	85.31	–	25.36	9,176	18.6%
WS & WA	713	87.56	88.56	36.93	6,625	19.1%
Advertiser	136	87.50	88.15	35.93	22,496	100.0%
Non-advertiser	577	87.58	88.65	37.16	3,440	0.0%
Advertiser	516	86.14	–	26.76	29,853	100.0%
Non-advertiser	2,237	85.84	–	28.72	3,768	0.0%
Panel B: All Wines	10,935	85.39	–	31.85	7,094	9.8%
Advertiser	1,071	85.52	–	29.01	27,459	100.0%
Non-advertiser	9,864	85.37	–	32.17	4,726	0.0%
Best Buy	126	86.10	–	9.35	49,573	26.2%
Spectator Selection	57	90.91	–	25.02	13,203	36.8%
Highly Recommended	104	93.35	–	51.42	3,293	15.4%
Cellar Selection	66	94.58	–	125.08	4,743	16.7%
Panel C: Winery-level Observations						
U.S. Wineries						
Advertiser	76	84.87	–	20.89	196,408	93.3%
Non-advertiser	735	85.70	–	28.47	10,964	0.0%
All Wineries						
Advertiser	183	84.91	–	27.54	134,170	91.8%
Non-advertiser	3,153	85.14	–	26.09	10,466	0.0%

Panel A is restricted to U.S. wines reviewed by *Wine Spectator* in 2000; Panel B includes wines produced in other countries. Wines reviewed from barrel or as part of special tastings are excluded. Observation are classified as coming from an advertiser if the winery advertised at least once in the prior twelve months. Column "% Advert." contains the percentage of wines produced by advertisers. Panel C contains winery-level statistics estimated from the characteristics of the wines reviewed by *Wine Spectator*.

advertised in *Wine Spectator* prior to being reviewed. Moreover, it indicates that wines produced by advertisers have slightly higher *Wine Spectator* ratings (86.14 versus 85.84), slightly lower prices ($26.76 versus $28.72), and are made in much larger quantities (29,853 cases versus 3,768 cases) than wines produced by non-advertisers. The higher ratings (for a given price and production level) are consistent with biased wine ratings but are also consistent with more advertising by higher-quality wineries. Comparing wine for which *Wine Advocate* ratings are available to those for which they are not, reveals that wines reviewed by *Wine Advocate* are more expensive and produced in smaller quantities than the average wine reviewed by *Wine Spectator.*[14] The higher average *Wine Spectator* rating for wines with *Wine Advocate* ratings reflects the fact that *Wine Advocate* rarely publishes ratings of

[14] I did not collect data on the set of wines reviewed by *Wine Advocate* but not reviewed by *Wine Spectator*. Therefore, if *Wine Advocate* tends to review smaller–production wines than does *Wine Spectator*, the differences in panel A understate the true differences between the two publications.

less than 85 points and that the ratings in the two publications are positively correlated. The average differences in the *Wine Spectator* and *Wine Advocate* ratings of advertisers and non-advertisers suggest a difference-in-difference estimated advertising bias of 0.42 points (within the nonrandom sample of wines reviewed by both publications).

Panel B focuses on the full set of wines reviewed by *Wine Spectator* in 2000. The comparison of wines produced by advertisers and non-advertisers is qualitatively similar to the comparison in Panel A. *Wine Spectator* awards, collectively bestowed upon approximately 3.2 percent of the wines reviewed each year, are of four types: *Best Buys* are "wines of value [with] solid scores, modest prices, wide distribution"; *Cellar Selections* are "wines we believe will improve most from additional bottle age and show the greatest potential as collectibles"; *Highly Recommended* wines are "noteworthy wines selected from among the highest-scoring wines in the issue"; and *Spectator Selections* are "*[Wine Spectator's]* highest recommendations in each issue.... they are the wines we think would make the best purchases based on a combination of rating, price, and availability." Average ratings, prices, and production levels differ across the four *Wine Spectator* awards in ways consistent with the stated criteria. Although I do not test for biased product coverage, the fact that only 9.8 percent of wines come from wineries that advertise suggests that wineries need not advertise in order to have their products reviewed. The fact that 18.7 of U.S. wines come from advertisers, indicates that foreign wineries are less likely to advertise in *Wine Spectator* than domestic wines (although this aggregate percentage masks substantial cross country differences).

Finally, Panel C compares wineries that advertise at any point in 2000 to non-advertisers. For each winery, I calculated production-weighted average *Wine Spectator* ratings and prices, total production, and the percentage of wines reviewed in 2000 when classified as an advertiser (*i.e.*, for which "Ad Dummy" equals one). I then averaged the winery-level statistics across the two sets of wineries. Of the 811 U.S. wineries whose wines were reviewed by *Wine Spectator*, only 76 advertised in *Wine Spectator* in 2000. Moreover, the estimated average production level of these 76 advertisers is approximately 18 times greater than the average production level of non-advertisers. The difference in production levels is smaller, but still substantial, when including foreign wineries. Also, for both sets of wineries, the production-weighted average *Wine Spectator* ratings of advertisers are less than those of non-advertisers. Comparing the number of wines reviewed by *Wine Spectator* to the number of wineries, indicates that each year *Wine Spectator* reviews, on average, approximately three wines per winery.

III. Empirical Framework

A. *Uni-Dimensional Quality*

Let subscript i represent a particular wine produced by winery j from grapes grown in region k during vintage l. For example, when wine i is the 1995 Beringer Private Reserve

Napa Valley Cabernet Sauvignon, *j*, *k*, and *l* are Beringer, Napa Valley, and 1995, respectively. Assume that quality can be measured along a single dimension and let Q_i^* denote the quality of wine *i*. Because Q_i^* is unobservable, wine ratings are intended to provide consumers with estimates of Q_i^*. Let WS_{it} be the rating awarded to wine *i* by *Wine Spectator* on date *t* and AD_{jt} be a stock measure of winery *j*'s advertising in *Wine Spectator* through (but not including) date *t*. Assume that *Wine Spectator* observes Q_i^* with noise and that WS_{it} is linear and strictly increasing in this (noisy) measure of quality.[15] Formally,

$$\begin{aligned} WS_{it} &= \alpha + \beta AD_{jt} + \gamma(Q_i^* + \tilde{\varepsilon}_{it}) \\ &\equiv \alpha + \beta AD_{jt} + \gamma Q_i^* + \varepsilon_{it} \end{aligned} \tag{1}$$

where ε_{it} is a random variable with zero mean (so that $Q_i^* + \tilde{\varepsilon}_{it}$ is an unbiased estimate of Q_i^*) and WS_{it} is a function of the advertising measure AD_{jt}. Under the null hypothesis of unbiased ratings, β equals zero, whereas under the alternative hypothesis of (positively) biased ratings, β is positive.

If Q^* were observed by the econometrician, it would be straightforward to test whether the coefficient on *AD* were significantly different from zero. That Q^* is unobserved introduces econometric difficulties. Specifically, unless Q^* and *AD* are uncorrelated, a regression of WS on *AD* alone suffers from an omitted-variables bias. In such a regression, the coefficient on *AD* will be biased upward if wineries that produce higher quality wines are more likely to advertise in *Wine Spectator* (*i.e.*, quality and advertising are strategic complements), and it will be biased downward if wineries that produce lower quality wines are more likely to advertise in *Wine Spectator* (*i.e.*, quality and advertising are strategic substitutes).

There are two ways to control for unobserved quality. First, I can model Q^* as a function of observable characteristics and use those characteristics as additional explanatory variables in a regression of *WS* on *AD*. For this approach to remove the omitted-variables bias, it must control for the variation in Q^* that is correlated with *AD*. Consider a model which decomposes quality into three components:

$$Q_i^* = \theta_i + \eta_j + \eta_{kl} \tag{2}$$

where θ_i depends on actions taken by the winery and winemaker during vintage *k* and includes a shock specific to wine *i*, η_j is a winery fixed effect, and η_{kl} is a region-by-vintage fixed effect.

[15] For convenience, I treat WS_{it} as a continuous variable rather than a variable that can only take on integers between 50 and 100. Simulations suggest that rounding from continuous variables to integers introduces negligible measurement error in this context. Because ratings of 50 and 100 are rare, estimation techniques that treat these ratings as censored yield results that are virtually identical to those reported in the text.

The region and vintage fixed effect acknowledges that fluctuations in quality across vintages and regions are driven by shared fluctuations in the weather. Since high quality grapes are a necessary input into the production of high quality wine, weather patterns during the vintage are an important determinant of final quality.[16] Moreover, because my tests focus on releases of U.S. wines within a single year, the number of region-by-vintage fixed effects is relatively small, allowing me to control for this source of variation in Q^*. The winery fixed effect acknowledges that some wineries consistently possess better vineyards (within region l) and employ better winemakers than do other wineries. However, because my sample consists of wines reviewed in a single year and because a large number of wineries only have one or two wines reviewed each year, I cannot control for variation in η_j.[17] It is also not possible to control directly for θ_i although I can include production level and price of wine i as proxies. The expected relationship between θ_i and production level is negative because growing fewer grapes per acre produces higher quality grapes.[18] To the extent that higher prices reflect production costs associated with producing higher quality wine, the expected relationship between θ_i and price is positive.

The second way I can control for unobserved quality is to use the published rating for wine i from *Wine Advocate* as a proxy for Q_i^*. Since the *Wine Advocate* rating is based on a tasting of wine i, it is likely to capture variation in quality across regions and vintages (η_{kl}) as well as how wine i compares to wines from the same region and vintage ($\theta_i + \eta_j$). And because *Wine Advocate* does not accept advertising, *Wine Advocate* ratings should be unbiased estimates of Q_i^*. For these reasons, *Wine Advocate* ratings may be better proxies for Q_i^* than a model based on fixed effects and observable characteristics.

However, a complication arises from the fact that I only observe *Wine Advocate* ratings for a subset of the wines for which I observe *Wine Spectator* ratings. Let WA_{it}^* be the rating assigned to wine i by *Wine Advocate* on date t.[19] Assume that

$$WA_{it}^* \equiv Q_i^* + \tilde{v}_{it} \tag{3}$$

[16] Ashenfelter, Ashmore, and Lalonde (1995) demonstrate that 83 percent of across-vintage variation in the auction prices of mature Bordeaux can be explained by winery fixed effects, years since vintage, and variables that summarize weather patterns during the vintage. Ashenfelter (2008) uses the same model to explain an even greater percentage of the variation in more recent data. Byron and Ashenfelter (1995) perform a similar analysis for Penfolds Grange. In all three studies, auction prices, rather than expert opinion, are taken as measures of quality. Studying the quality of wines from the Mosel Valley in Germany, Ashenfelter and Storchmann (2009) extend measures of quality to include both retail prices and auction prices. In addition, they augment the traditional hedonic model with one based on solar radiation.

[17] Put differently, adding winery fixed effects uses the variation in advertising within each winery to identify biased ratings. However, because I possess a short time series, and advertising is highly persistent among U.S. wineries, there is little within–winery variation in advertising within my sample.

[18] For example, depending on how much effort is spent reducing yields, they can range from less than 2 tons per acre to 8 or more tons per acre.

[19] For notational convenience, I assume that *Wine Spectator* and *Wine Advocate* both review wine i at time t. Since I am interested in the influence of advertising in *Wine Spectator* through (but not including) time t on *Wine Spectator* reviews published at time t, the assumption is largely without cost. In practice, *Wine Advocate* tends to publish the rating for any particular wine before *Wine Spectator* does.

where $\tilde{v}_{it}$ is a random variable with zero mean that is uncorrelated with $\tilde{\varepsilon}_{it}$ (the noise in the *Wine Spectator* rating).[20] There are two reasons that I might observe WS_{it} and not observe WA_{it}^{*} First, *Wine Advocate* may not have tasted wine *i*. Second, because *Wine Advocate* rarely publishes ratings for wines rating less than 85 points, *Wine Advocate* may have tasted wine *i* and rated it below 85 points. Combining these conditions, the published wine rating WA_{it} is given by

$$WA_{it} = \begin{cases} WA_{it}^{*} & \text{if } WA_{it}^{*} \geq 85 \\ \text{missing} & \text{if } WA_{it}^{*} < 85 \text{ } or \text{ wine } i \text{ was not tasted.} \end{cases} \tag{4}$$

That I do not know whether WA_{it} is missing because wine *i* was not]tasted or because wine *i* was tasted and received a rating less than 85, differentiates equation (4) from a standard censored regression problem. Instead, one can think of equation (4) as arising from a model in which *Wine Advocate* tastes every wine that *Wine Spectator* tastes but applies different censoring criteria to different wines. Some wines are coded as missing when WA_{it}^{*} is less than 85, other wines are coded as missing for all values of WA_{it}^{*}, and I lack data on which censoring point applies to any particular wine. As noted in Amemiya (1985, p. 363), models in which censoring points vary unobservably across observations typically cannot be estimated. To conduct tests of biased wine ratings several additional assumptions are required.

First, I assume that $\tilde{\varepsilon}_{it}$ and $\tilde{v}_{it}$ are jointly normally distributed. Under this assumption, equations (1) and (3) form a (nonstandard) simultaneous-equations Tobit model. Rather than estimate a model with WS_{it} as the dependent variable and (censored) WA_{it} and AD_{jt} as independent variables, I treat WA_{it} as a censored dependent variable and WS_{it} and AD_{jt} as independent variables. Solving equation (1) for Q_{i}^{*} and plugging into equation (3) yields

$$\begin{aligned} WA_{it}^{*} &= -\frac{\alpha}{\gamma} - \frac{\beta}{\gamma} AD_{jt} + \frac{1}{\gamma} WS_{it} - \frac{1}{\gamma}\varepsilon_{it} + \tilde{v}_{it} \\ &\equiv \overline{\alpha} + \overline{\beta} AD_{jt} + \overline{\gamma} WS_{it} + \overline{v}_{it} \end{aligned} \tag{5}$$

where WS_{it} and $\overline{v}_{it}$ are correlated because both terms contain ε_{it}. Plugging equation (5) into equation (4) yields

$$WA_{it} = \begin{cases} WA_{it}^{*} & \text{if } \overline{\alpha} + \overline{\beta} AD_{jt} + \overline{\gamma} WS_{it} + \overline{v}_{it} \geq 85 \\ \text{missing} & \text{if } \overline{\alpha} + \overline{\beta} AD_{jt} + \overline{\gamma} WS_{it} + \overline{v}_{it} < 85 \text{ } or \text{ wine } i \text{ was not tasted.} \end{cases} \tag{6}$$

Equation (6) is the reduced-form for equations (1) and (3). The test that β is positive in equation (1) can be recast as a joint test of whether $\overline{\beta}$ is negative and $\overline{\gamma}$ is positive. Of course, once

[20] The assumption that WA_{it}^{*} equals $Q_{i}^{*} + \tilde{v}_{it}$, rather than $\alpha_{wa} + \gamma_{wa} Q_{i}^{*} + \tilde{v}_{it}$, is without loss of generality. Both specifications give rise to the same equation (5).

$\bar{\beta}$ and $\bar{\gamma}$ have been estimated, the structural parameter β can be estimated as the negative of the ratio of $\bar{\beta}$ to $\bar{\gamma}$ (and its standard error can be estimated using the delta method).

Second, because WS_{it} is correlated with ε_{it}, equation (6) suffers from endogeneity. To eliminate the associated bias, I predict WS_{it} using instrumental variables which are correlated with Q_i^* but likely to be uncorrelated with the idiosyncratic tasting shock ε_{it}. Candidates for instruments include price and production. I then replace the actual *Wine Spectator* rating in equation (6) with the predicted *Wine Spectator* rating.[21]

Finally, to overcome the problem of unobserved censoring points, I assume that wine i is tasted by *Wine Advocate* with fixed probability p_i. In other words, I assume that a *known* fraction of the observations are missing at random and the remainder are missing because WA_{it}^* is less than 85.[22] To estimate equation (6), I assume a particular value of p and then estimate the log likelihood function conditional on this p. (Reasonable values of p are discussed in Section 4.)

Let o_i equal one if WA_{it}^* is observed and zero otherwise. The likelihood function corresponding to equation (6) is given by

$$L = \prod \left(\sigma^{-1} \phi(\cdot)\, p \right)^{o_i} \left(1 - p\left(1 - \Phi(\cdot)\right)\right)^{1-o_i} \tag{7}$$

where σ is the standard deviation of $\bar{v}_{it}$, ϕ and Φ correspond to the pdf and cdf of a standard normal random variable, and

$$\phi(\cdot) = \phi((WA_{it} - \bar{\alpha} - \bar{\beta} AD_{jt} - \bar{\gamma} WS_{it})/\sigma)$$
$$\Phi(\cdot) = \Phi((85 - \bar{\alpha} - \bar{\beta} AD_{jt} - \bar{\gamma} WS_{it})/\sigma).$$

The likelihood function in equation (7) differs from the standard Tobit likelihood function in that it contains p, the probability that *Wine Advocate* tasted wine i. The log likelihood function is given by

$$\ln L = \sum o_i (\ln \sigma^{-1} + \ln \phi(\cdot)) + \sum o_i \ln p + \sum (1 - o_i) \ln (1 - p(1 - \Phi(\cdot))) \tag{8}$$

$$\ln L' = \sum o_i \, (\ln \sigma^{-1} + \ln \phi(\cdot)) + \sum (1 - o_i) \ln (1 - p(1 - \Phi(\cdot))) \tag{9}$$

[21] Nelson and Olson (1978) argue that replacing endogenous regressors with predicted values and maximizing a standard Tobit likelihood function yields consistent and asymptotically normal estimates of the Tobit coefficients. However, the standard errors implied by the information matrix understate the true standard errors because they ignore the prediction error. Moreover, Newey (1987) demonstrates that this technique does not generally yield efficient estimates. While the efficient estimator proposed in Newey (1987) cannot readily be extended to the modified likelihood function that I estimate, to adjust for prediction error, I bootstrap my standard errors.

[22] Alternatively, I could assume a relationship between observable characteristics and the probability that wine i is tasted. However, I lack the data needed to estimate such a model.

where the summation involving ln p in equation (8) can be dropped because once the value of p is fixed, it is a constant. Coefficients are estimated by maximizing the likelihood function given in equation (9). Note that because the actual *Wine Spectator* rating is replaced with the predicted *Wine Spectator* rating, the standard errors estimated via maximum likelihood understate the true standard errors. To adjust for the use of predicted regressors, I estimate standard errors via bootstrapping.

B. *Multi-Dimensional Quality*

The discussion above assumes that quality can be measured along a single dimension. Under this assumption, WA_{it} is a good proxy for Q_i^* and I can interpret a positive coefficient on *AD* as evidence of biased wine ratings. Interpreting the coefficient on *AD* becomes more difficult, however, if quality is multi-dimensional, the two publications favor different dimensions, *and* these differences are correlated with wineries decisions about whether and how much to advertise.

Let Q_i^{ws} be the quality of wine i as perceived by *Wine Spectator* and let Q_i^{wa} be the quality of wine i as perceived by *Wine Advocate*. Equation (1) becomes

$$\begin{aligned} WS_{it} &= \alpha + \beta AD_{jt} + \gamma Q_i^{ws} + \varepsilon_{it} \\ &= \alpha + \beta AD_{jt} + \gamma Q_i^{wa} + \gamma(Q_i^{ws} - Q_i^{wa}) + \varepsilon_{it} \end{aligned} \tag{10}$$

Using WA_{it} to proxy for unobserved quality does not control for unobserved differences between Q_i^{ws} and Q_i^{wa}. If these differences are correlated with wineries' decisions about how much to advertise in *Wine Spectator*, regressions of *WS* on *AD* and *WA* suffer from the sort of omitted-variables bias discussed above. For example, if wineries that expect relatively higher *Wine Spectator* ratings are more likely to advertise in *Wine Spectator*, the estimated coefficient on *AD* will be positive even if ratings are unbiased and the true coefficient on *AD* is zero.

Empirically, there are two ways to address this type of omitted-variables bias. First, I can use measures of advertising intensity that are less sensitive to a potential correlation between AD_{jt} and the difference between Q_i^{ws} and Q_i^{wa}. For example, rather than measure AD_{jt} as the number of pages of advertising taken out by winery j through (but not including) time t, I could measure AD_{jt} as a dummy variable that equals one if and only if winery j has previous advertised in *Wine Spectator*. Because *Wine Advocate* does not accept advertising and *Wine Spectator* has the largest circulation of any U.S. wine publications, U.S. wineries that do not produce "*Wine Advocate* style" wines may still advertise in *Wine Spectator*, just less often. Second, I can attempt to classify wineries as making either "*Wine Spectator* style" wines or "*Wine Advocate* style" wines and ask whether the impact of *AD* on ratings is differentially higher for the "*Wine Spectator* style" wineries. Data on which wineries cite *Wine Spectator* ratings and awards in their advertisements provide me with an (imperfect) way to identify "*Wine Spectator* style" wineries.

IV. Empirical Results

A. *Testing for Biased Ratings*

In this section, I test whether *Wine Spectator* biases its ratings to favor advertisers. Each test for biased ratings is run using two measures of advertising. The first measure is a dummy variable that equals one if the wine comes from a winery that advertised in *Wine Spectator* in the 12 months prior to the issue published on date *t*. It allows me to ask whether advertisers are treated differently on average than non-advertisers. The second measure is a function of the number of pages of advertising that winery placed in *Wine Spectator* in the 12 months prior to time *t*. (The function is the natural logarithm of one plus the number of pages of advertising, which ranges from 0 to approximately 2.5.) It allows me to ask whether bias is increasing in the amount that wineries advertise.

These measures differ in two ways from the measures that *Wine Spectator* would likely use to determine advertising status. First, my advertising dummy classifies a winery as an advertisers in the issue after its first advertisement, but does not stop classifying the winery as an advertiser until 12 months after its last advertisement. Therefore, it is possible that only a subset of the wineries which I classify as advertisers are actively advertising in *Wine Spectator*. Second, while *Wine Spectator* would presumably rank wineries based on expected advertising revenue, I focus on pages of advertising placed in *Wine Spectator* prior to the review of each wine. This measure ignores the facts that some pages of the magazine are more expensive to advertise in than others and that *Wine Spectator* derives additional advertising revenue from *Wine Spectator Online*. Also, if a winery undertakes a large advertising campaign, *Wine Spectator* might immediately classify the winery as a large advertiser whereas my page measure will initially classify it as a relatively small advertiser. However, adjusting the advertising measures to include future advertising will bias the tests towards finding evidence of biased wine ratings if wineries that receive favorable reviews at time *t* choose to advertise those reviews at time *t +1*. Moreover, winery advertising is highly persistent in my sample. Wineries that advertise in 1999 are quite likely to advertise throughout 2000, and there are only a few wineries that advertised in 2000 and did not advertise in 1999. Overall, I expect that my advertising measures are reasonable proxies for which wineries advertise and how much they advertise.

Table 4 contains my initial tests for biased wine ratings. These tests use the wine ratings and advertising data collected from *Wine Spectator*, but ignore wine ratings from *Wine Advocate*. Within this sample, I find little evidence that advertising is correlated with ratings. In columns (1) and (4), I present univariate regressions of *Wine Spectator* ratings on each of the advertising measures. (All of the standard errors in Table 4 are clustered on winery.) While both estimated coefficients are positive, only the coefficient on advertising pages is statistically significant from zero (and only at the 10-percent level). In addition, the implied impact of advertising on ratings is modest. For example, an increase from 0 to 1.286 (the average (natural log) page measure for wineries that advertise) is predicted to increase the *Wine Spectator* rating by 0.406 points.

Table 4
OLS-Based Tests for Biased Wine Ratings

Dependent Variable:	Wine Spectator *Rating*					
Advertising Measure:		*Dummy*			*Ln Pages*	
	(1)	(2)	(3)	(4)	(5)	(6)
Lagged Advertising	0.305 (0.297)	0.318 (0.272)	–0.175 (0.213)	0.316* (0.191)	0.367** (0.192)	0.096 (0.152)
Ln Production		–0.320*** (0.055)	0.317*** (0.052)		–0.339*** (0.054)	0.285*** (0.052)
Ln Price			4.285*** (0.179)			4.260*** (0.179)
Fixed Effects:						
Varietal	No	Yes***	Yes***	No	Yes***	Yes***
Region × Vintage	No	Yes***	Yes***	No	Yes***	Yes***
Sample Size	2,753	2,677	2,668	2,753	2,677	2,668
Adjusted R^2	0.0009	0.2838	0.4433	0.0020	0.2851	0.4432

The table reports estimated coefficients from OLS regressions restricted to the sample of U.S. wines reviewed in *Wine Spectator* in 2000. The dependent variable is the *Wine Spectator* rating of wine *i*. Independent variables include a lagged advertising measure, natural logarithm of price, natural logarithm of production, varietal fixed effects, and region-by-vintage fixed effects. Columns (1) through (3) measure advertising with a dummy variable that equals one if the winery advertised in *Wine Spectator* at least once in the prior 12 months. Columns (4) through (6) measure advertising as (the natural log of one plus) the number of pages of advertising taken out by the winery in the prior 12 months. Standard errors are clustered on winery. Coefficients (and fixed effects) significantly different from zero (in two-sided tests) at the 10-percent level are marked with *, 5-percent level are marked with **, and 1-percent level are marked with ***.

Of course, to the extent that winery decisions about how much to advertise in *Wine Spectator* depend on the quality of their wines, the univariate regressions suffer from omitted-variables bias. In an attempt to control for the unobserved quality of each wine, columns (2) and (4) include the natural logarithm of production, varietal fixed effects, and region-by-vintage fixed effects as additional regressors. Interestingly, while the estimated coefficients on these controls are statistically significant, the estimated coefficients on the advertising measures are essentially unchanged. However, when I add the natural logarithm of price as an additional control for quality, I find no evidence that advertising is positively correlated with *Wine Spectator* ratings.

To better control for quality, Table 5 presents tests for biased wine ratings based upon the modified Tobit framework derived above. This framework uses *Wine Advocate* ratings to control for the unobserved quality of each wine reviewed by *Wine Spectator.* However, because *Wine Advocate* ratings are censored, I treat them as the dependent variable and treat *Wine Spectator* ratings and the advertising measures as independent variables. The tests for biased wine ratings can then be recast as a joint test that the coefficient on *Wine Spectator* ratings is positive and the coefficient on the advertising measure is negative (so that *Wine Advocate* ratings are consistently lower than *Wine Spectator* ratings for wines from advertisers). Alternatively, the coefficients obtained by maximizing the modified Tobit likelihood function can be used to recover direct estimates of how each advertising

Table 5
Tobit-Based Tests for Biased Wine Ratings

Dependent Variable:	Wine Advocate *Rating*							
Advertising Measure:	*Dummy*				*Ln Pages*			
WS *Rating:*	*Actual*		*Predicted*		*Actual*		*Predicted*	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Lagged Advertising	–0.353	–1.076**	–1.704***	–1.045**	–0.102	–0.736*	–1.319***	–0.915*
	(0.365)	(0.551)	(0.611)	(0.575)	(0.398)	(0.481)	(0.501)	(0.469)
Lagged Ad Ratio			2.301*				2.664**	
			(1.225)				(1.237)	
Retasted Dummy				2.474**				1.614
				(1.237)				(1.086)
Lagged Ad				–1.027				1.538
× Retasted Dummy				(3.206)				(1.854)
WS Rating	0.620***	1.115***	1.109***	1.094***	0.620	1.117***	1.106***	1.090***
	(0.149)	(0.439)	(0.283)	(0.397)	(0.427)	(0.374)	(0.309)	(0.351)
Sample Size	2,753	2,673	2,673	2,673	2,753	2,673	2,673	2,673
	Direct Effect of Independent Variable on *Wine Spectator* Rating							
Lagged Advertising	0.569	0.964**	1.536***	0.955**	0.165	0.659*	1.193***	0.840**
	(0.627)	(0.408)	(0.509)	(0.474)	(0.307)	(0.338)	(0.431)	(0.340)
Lagged Ad Ratio			–2.076*				–2.409**	
			(1.073)				(1.132)	
Retasted Dummy				–2.260**				–1.481
				(0.989)				(0.940)
Lagged Ad				0.938				–1.411
× Retasted Dummy				(2.098)				(1.514)

The top panel reports coefficients from the Modified Tobit described in Section 3, estimated using the full sample of U.S. wines reviewed by *Wine Spectator* in 2000. The dependent variable equals the *Wine Advocate* rating when the rating is observed and greater than or equal to 85, and it equals zero otherwise. Independent variables include a lagged advertising measure and *Wine Spectator* ratings. Natural logarithm of production and price used as instruments when predicting *Wine Spectator* rating. Columns (1) through (3) measure advertising with a dummy variable that equals one if the winery advertised at least once in the prior 12 months. Columns (4) through (6) measure advertising as (the natural log of one plus) the number of pages of advertising taken out by the winery in the prior 12 months. Because six of the specifications involve predicted regressors, standard errors in the top panel are estimated via bootstrapping (with 1000 replications). The bottom panel reports the direct effect of each independent variable on *Wine Spectator* ratings. For example, the direct effect of advertising measure on *Wine Spectator* rating estimated as the negative of the the coefficienton lagged advertising measure divided by the coefficient on *Wine Spectator* rating; the magnitude is directly comparable to estimates in Table 4. Standard errors for the direct effects are estimated via the delta method. Coefficients (and fixed effects) significantly different from zero (in two-sided tests) at the 10-percent level are marked with *, 5-percent level are marked with **, and 1-percent level are marked with ***.

measure influences *Wine Spectator* ratings.[23] The top panel of Table 5 reports maximum likelihood estimates based on the modified Tobit maximum likelihood function. Because several of the specifications include predicted ratings, standard errors are estimated via bootstrapping (based on 1000 replications). The bottom panel of Table 5 reports the direct effect of each independent variable on the *Wine Spectator* rating, with standard errors calculated using the delta method.

The coefficients estimated via maximum likelihood are functions of the probability that *Wine Advocate* tasted wines for which it did not publish reviews. Approximately 40 percent of the *Wine Spectator* ratings of U.S. wines are below 85 points. Assuming this holds true in the distribution of *Wine Advocate* ratings as well, the 680 *Wine Advocate* ratings at and above 85 points imply approximately 1,133 total reviews. As a ratio to the total number of wines reviewed by *Wine Spectator*, this implies that the probability of *Wine Advocate* having tasted any particular bottle of wine is approximately 0.41. Therefore, I set p equal to 0.41.[24]

The results in columns (1) and (5) are based on specifications that include actual *Wine Spectator* ratings with each lagged advertising measure. The estimated Tobit coefficients on actual *Wine Spectator* ratings are both 0.620, but the estimate in column (5) is statistically indistinguishable from zero. In addition, the estimated Tobit coefficients on both advertising measures (and the implied direct effects of those measures on *Wine Spectator* ratings) are statistically indistinguishable from zero. In other words, these specifications yield no evidence of bias. However, because actual *Wine Spectator* ratings are assumed to be noisy measures of true quality, these initial specifications suffer from errors-in-variables.

The remaining six columns of Table 5 replace actual *Wine Spectator* ratings with predicted *Wine Spectator* ratings, where the natural logarithm of production and price are used as instruments in first stage regressions. In each case, the estimated Tobit coefficient on *Wine Spectator* ratings goes from 0.620 to slightly above one.[25] Correcting for the measurement error in *Wine Spectator* ratings, the estimated coefficients in columns (2) and (6) are consistent with *Wine Spectator* bestowing slightly higher ratings on advertisers. The estimated Tobit coefficients in column (2) imply that wines from advertisers receive 0.964 more points from *Wine Spectator* than comparable wines from non-advertisers (significant at the 5-percent level). The results in column (6) imply that the level of bias is increasing

[23] For example, the effect of an advertising measure on *Wine Spectator* ratings—β in equation (1) —can be estimated as the negative of the ratio of the Tobit coefficient on the advertising measure to the Tobit coefficient on *Wine Spectator* ratings. This estimate is directly comparable to estimates in Table 4.

[24] I censor the *Wine Advocate* ratings for the 33 wines with published *Wine Advocate* ratings less than 85 points. When I re-estimate the specifications in Table 5 for values of p between 0.3 and 0.5, the results are qualitatively similar to those reported in the text, except that estimated bias is slightly higher when p equals 0.3 and slightly lower when it equals 0.5.

[25] Ignoring the fact that *Wine Advocate* is censored, if *Wine Advocate* and *Wine Spectator* ratings are noisy measurements of the same true quality level Q_i^*, univariate OLS regressions of *Wine Advocate* ratings on *Wine Spectator* suffer from errors-in-variables and the coefficient on *Wine Spectator* is biased towards zero. In this case, using instrumental variables to remove the measurement error should produce an estimated coefficient on the *Wine Spectator* rating near one.

in the level of prior advertising, with the average advertiser receiving 0.847 more points (0.659 times 1.286) than a non-advertiser.

Under the identifying assumption that *Wine Advocate* and *Wine Spectator* rate wines using the same standards, the results in columns (2) and (6) constitute evidence of biased wine ratings. Columns (3) and (7) attempt to test this identifying assumption. Lagged advertising ratio is the ratio of the number of pages citing *Wine Spectator* ratings and awards to the total number of pages of advertising placed by each winery. If some wineries produce "*Wine Spectator* style" wines and these wineries are more likely to cite *Wine Spectator* in their advertisements, lagged advertising ratio should be highest for the "*Wine Spectator* style" wineries and the Tobit coefficients on lagged advertising ratio should be negative.[26] Instead, the estimated Tobit coefficients on lagged advertising ratio are positive and the estimated Tobit coefficients on the other lagged advertising measures are more negative than in columns (2) and (6). For example, according to the estimates in the bottom panel of column (2), wines from advertisers for which lagged advertising ratio equals zero receive 1.536 more points than comparable wines from non-advertisers while advertisers for which lagged advertising ratio equals one receive −0.540 less points (1.536 plus –2.076) than comparable wines from non-advertisers (although this difference is not statistically different from zero). Another way to ask whether the biased ratings results are being driven by "*Wine Spectator* style" wineries is to compare the estimates using the advertising dummy to the estimates using pages of advertising. If "*Wine Spectator* style" wineries advertise more often than other wineries, this might explain why the amount of bias is increasing in pages of advertising. However, the bias implied by both advertising measures is quite similar for the average advertiser.

B. *Testing Whether Biased Ratings Arise from Selective Retasting*

Although *Wine Spectator's* use of blind tastings would appear to rule out biased ratings, biased ratings potentially could arise through selective retasting. Specifically, if *Wine Spectator* were more likely to retaste wines from wineries that advertise, and if published scores tend to represent the best of several bottles tasted, these facts would bias *Wine Spectator's* published wine ratings in favor of advertisers. Let Retaste_{it} be a dummy variable that equals one if *Wine Spectator* reports tasting wine *i* more than once before publishing its rating at time *t*, and zero otherwise. To test whether the observed bias in wine ratings can be explained by selective retastings on the part of *Wine Spectator*, in columns(4) and (8), I include $Retaste_{it}$ and $Retaste_{it}$ interacted with AD_{jt} as additional explanatory variables. Within the full sample of U.S. wines, I find little evidence that selective retasting can explain the higher ratings earned by *Wine Spectator's* advertisers. First, I continue to find evidence that advertisers receive an additional point from *Wine Spectator*. Second, the coefficients on the interaction terms provide no support for the hypothesis that retastings benefit advertisers differentially more than non-advertisers.

[26] I do not know whether advertisements that cite *Wine Spectator* ratings and awards also cite other publications. To the extent that they do, the advertising ratio may reflect the extent to which wineries attempt to establish quality through favorable reviews rather than signal that their wine is made in a "*Wine Spectator* style."

In a final attempt to determine the mechanism through which the positive (partial) correlation between advertising measures and *Wine Spectator* ratings might arise, Table 6 restricts tests for selective retasting to the nonrandom sample of U.S. wines for which I observe ratings from both publications. If *Wine Advocate* ratings prompt *Wine Spectator* to retaste wines—perhaps because *Wine Spectator* is reluctant to publish lower ratings for advertisers than their wines earned in *Wine Advocate*—this is precisely the sample of wines in which selective retastings might bias ratings. As motivation for restricting attention to this sample, coefficients estimated via probit in columns (1) and (2) demonstrate that *Wine Spectator* is significantly more likely to retaste wines that have been reviewed by *Wine Advocate*. (The marginal effect is estimated to be 3 percentage points, which is quite large given that only 4.4 percent of wines are retasted.) These probit specifications also provide some evidence that *Wine Spectator* is more likely to retaste advertisers' wines. The marginal effects range from 1.250 percentage points (p-value of 0.084) to 1.858 percentage points (p-value of 0.131).

Table 6
Do Biased Ratings Arise Through Selective Retastings?

Estimation Technique:	*Probit*		*OLS*	
Dependent Variable:	*Retasted Dummy*		$WS_{it} - WA_{it}$	
Advertising Measure:	*Dummy*	*Ln Pages*	*Dummy*	*Ln Pages*
	(1)	(2)	(3)	(4)
WA Tasted	3.080***	3.086***		
	(1.149)	(1.147)		
Lagged Advertising	1.858	1.250*	0.300	0.118
	(1.362)	(0.729)	(0.578)	(0.380)
Retasted Dummy			–2.699***	–1.952**
			(0.951)	(0.917)
Lagged Ad x Retasted			2.406	0.193
			(1.475)	(1.019)
Observed Pr (Retasted) × 100	4.403	4.403		
Predicted Pr (Retasted) × 100	4.217	4.208		
Sample Size	2,753	2,753	705	705
Pseudo or Adjusted R^2	0.0149	0.0157	0.0226	0.0148

Columns (1) and (3) measure advertising with a dummy variable that equals one if the winery advertised at least once in the prior 12 months. Columns (2) and (4) measure advertising as (the natural log of one plus) the number of pages of advertising taken out by the winery in the prior 12 months. Columns (1) and (2) estimate the probability that wine *i* is retasted by *Wine Spectator* as a function of a dummy variable that indicates whether *Wine Advocate* published a review for the wine and a lagged advertising measure. The probit specifications are estimated on the full set of U.S. wines reviewed by *Wine Spectator*. Marginal effects (multiplied by 100) are reported. Columns (3) and (4) regress the difference between the published *Wine Spectator* and *Wine Advocate* ratings on a lagged advertising measure, a dummy that indicates whether the wine was retasted by *Wine Spectator,* and the retasted dummy interacted with the lagged advertising measure. As such, columns (3) and (4) are restricted to the set of U.S. wines for which *Wine Advocate* published a rating. All standard errors are clustered on winery. Coefficients significantly different than zero at the 10-percent level are marked with *, 5-percent level are marked with **, and 1-percent level are marked with ***.

Columns (3) and (4) then regress differences between *Wine Spectator* and *Wine Advocate* ratings on a lagged advertising measure, a retasted dummy, and the retasted dummy interacted with the lagged advertising measure. In these specifications, the estimated coefficients on the lagged advertising measures are much lower than those implied by Table 5 and are statistically indistinguishable from zero. The negative coefficients on the retasted dummy indicate that wines which are retasted by *Wine Spectator* are rated, on average, between 1.95 and 2.70 points lower in *Wine Spectator* than in *Wine Advocate*. In column (3), the estimated coefficient on the retasted dummy interacted with the lagged advertising dummy is 2.406 (p-value of 0.104), suggesting that the lower ratings associated with retasting are limited to wines from non-advertisers. In other words, within this nonrandom sample, column (3) suggests that advertisers benefit more from retastings than do non-advertisers. Multiplying the 2.406 point bias by the 31 advertiser wines that were retasted by *Wine Spectator* and dividing by the 135 wines produced by advertisers (within this sample of 705 wines) yields an average pro-advertiser bias of 0.55 points—approximately half the level of bias implied by the estimates in Table 5. However, the differential retasting result does not hold in column (4), when the advertising dummy is replaced with pages of advertising. Therefore, while the selective retasting of advertisers' wines may contribute to the slightly higher ratings *Wine Spectator* bestows upon advertisers' wines, selective retasting is incapable of explaining a one point difference within either the full sample of U.S. wines or the subsample of wines for which I observe a *Wine Advocate* rating.

C. *Testing for Biased Awards*

The fact that *Wine Spectator* bestows awards upon a (small) fraction of the wines it reviews raises the possibility that awards are tied to advertising. I test for biased awards using data on the full sample of wines reviewed by *Wine Spectator* in 2000. In these tests, I ask whether *Wine Spectator* is more likely to bestow awards upon advertisers, controlling for the published *Wine Spectator* rating, price, and production level of each wine. Table 7 summarizes the characteristics of the wines that receive each *Wine Spectator* award. For all four awards, the percentage of awards going to advertisers (ranging from 15.4 percent to 36.8 percent) exceeds the percentage of wines reviewed from non-advertisers (9.8 percent). However, while the different awards are targeted at different types of wines, *Wine Spectator* has a stated preference for lower prices and higher production levels. Therefore, the higher percentage of awards going to advertisers may be explained by biased awards or by the lower average prices and higher average production levels of advertisers' wines.

Let Award_{it} be a dummy variable that equals one if wine i receives a particular award at date t and zero otherwise. Within the set of wines eligible for each type of award, the likelihood of receiving an award should be increasing in a wine's *Wine Spectator* rating and its production level, and decreasing in its price. If *Wine Spectator* is more likely to bestow awards on advertisers, the probability of receiving an award will also be an increasing function of advertising intensity:

$$\Pr(\text{Award}_{it} = 1) = h\,(WS_{it}, \text{price}_{it}, \text{production}_{i}, AD_{jt}) \tag{11}$$

Table 7
Characteristics of Wines Receiving *Wine Spectator* Awards, 2000

			WS *Rating*			*Price*			*Production*		
Award	*Obs.*	*% Advert.*	*Min.*	*Mean*	*Max.*	*Min.*	*Mean*	*Max.*	*Min.*	*Mean*	*Max.*
Best Buy	126	26.2%	84	86.10	90	$6	$9.35	$12	2,839	49,573	750,000
Spectator Selection	57	36.8%	89	90.91	95	9	25.02	45	2,500	13,203	70,000
Highly Recommended	104	15.4%	88	93.35	100	12	51.42	195	10	3,293	46,000
Cellar Selection	66	16.7%	90	94.58	100	35	125.08	750	80	4,743	27,000
No Award	10,582	9.4%	55	85.21	98	5	31.35	1,000	5	6,435	1,000,000
Sum	10,935										

This table compares the characteristics of wines that receive a *Wine Spectator* awards in 2000 to all of the other wines reviewed by *Wine Spectator* in 2000 (excluding wines tasted from barrel or as part of a special tasting). Observations are classified as coming from an advertiser if the winery advertised at least once in the prior twelve months. Column "% Advert." contains the percentage of wines produced by advertisers.

I estimate equation (11) via probit and test whether the coefficient on the advertising variable is positive and statistically different from zero (in two-sided tests). When interpreting the results, however, it is important to note that this test conditions on published *Wine Spectator* ratings. In the presence of biased wine ratings, equation (11) constitutes a test for additional advertising bias.[27]

The tests for biased awards focus on three sets of wine. Columns (1) and (2) test for biased awards using the full set of wines for which price and production are available. The dependent variable equals one if the wine received a *Wine Spectator* award and equals zero otherwise. Because different wines are eligible for different awards, the remaining columns focus on two subsets of wines. Columns (3) and (4) restrict attention to the set of wines priced $12 or less and produced in quantities of 2,800 cases or more. Within this set of wines, the dependent variable equals one if the wine received a *Best Buy* or *Spectator Selection* award and zero otherwise. In contrast, columns (5) and (6) restrict attention to wines rated 88 points or above, because these are the only wines that ever receive the *Cellar Selection, Highly Recommended*, and (remaining) *Spectator Selection* awards. Within this set of wines, the dependent variable equals one if the wine received any of these *Wine Spectator* awards and zero otherwise. Table 8 reports marginal effects (multipled by 100) for each of these specifications. Standard errors are clustered on winery.

[27] Assume, for example, that when deciding whether a wine should be designated *Best Buy* there is an optimal weight (based on the preferences of *Wine Spectator's* readers) to place on availability. The probit-based tests for biased awards ask whether *Wine Spectator* favors advertisers conditional on the observed weight *Wine Spectator* places on availability. In this framework it is not possible to test whether that weight is biased upward to favor the large-production wines produced by advertisers.

Table 8
Probit-Based Tests for Biased *Wine Spectator* Awards

Dependent Variable:	WS *Award*		*Best Buy*		*Other WS Award*	
Advertising Measure:	*Dummy*	*Ln Pages*	*Dummy*	*Ln Pages*	*Dummy*	*Ln Pages*
	(1)	(2)	(3)	(4)	(5)	(6)
Lagged Advertising	–0.019	–0.013	–0.517**	–0.334*	–0.836	–0.505
	(0.012)	(0.010)	(0.256)	(0.206)	(0.849)	(0.604)
WS Rating	0.070***	0.069***	1.124***	1.119***	2.905***	2.905***
	(0.020)	(0.020)	(0.291)	(0..293)	(0.279)	(0.280)
Ln Production	0.085***	0.084***	0.881***	0.871***	5.333***	5.315***
	(0.025)	(0.025)	(0.267)	(0.266)	(0.437)	(0.434)
Ln Price	–0.139***	–0.139***	–4.592***	–4.557***	–6.094***	–6.110***
	(0.039)	(0.039)	(1.281)	(1.296)	(0.826)	(0.826)
Observed Pr(Award) × 100	4.228	4.228	13.690	13.690	12.662	12.662
Predicted Pr(Award) × 100	0.046	0.046	0.722	0.716	4.415	4.413
Sample Size	7,828	7,828	935	935	2,614	2,614
Pseudo R^2	0.452	0.452	0.433	0.432	0.340	0.340

This table reports marginal effects (multiplied by 100) from probit regressions that predict whether wine *i* receives an award from *Wine Spectator* in 2000. "WS Award" equals one if the wine received one of Wine Spectator's four awards and zero otherwise. "Best Buy" is restricted to the set of wines costing $12 or less dollars and production is 2,800 cases or greater. It equals one if the wine received *Best Buy* or *Spectator Selection* and zero otherwise. "Other WS Award" is restricted to the set of wines scoring 88 points or above. It equals one if the wine received any of *Wine Spectator's* awards and zero otherwise. The observed probability of an award corresponds to the mean of the dependent variable; the predicted probability corresponds to the predicted probability of receiving an award, evaluated at the means of the independent variables. Columns (1), (3), and (5) measure advertising with a dummy variable that equals one if the winery advertised at least once in the prior 12 months. Columns (2), (4), and (6) measure advertising as (the natural log of one plus) the number of pages of advertising taken out by the winery in the prior 12 months. Standard errors are clustered on winery. Coefficients significantly different from zero (in two-sided tests) at the 10-percent level are marked with *, 5-percent level are marked with **, and 1-percent level are marked with ***.

Consistent with *Wine Spectator*'s stated criteria, the estimated probability of receiving an award is increasing in both the *Wine Spectator* rating and the natural logarithm of production and decreasing in the natural logarithm of price. In particular, increasing the *Wine Spectator* rating by one point is associated with a one percentage point increase in the probability of receiving the *Best Buy* award (relative to an observed probability of almost 14 percent) and a three percentage point increase in the probability of receiving any other award (relative to an observed probability of almost 13 percent). Therefore, biased ratings could translate into biased awards. However, the estimated coefficients on the advertising measures are negative, and in some cases statistically significant from zero. Therefore, despite the fact that *Wine Spectator* would appear to have more discretion over awards than ratings, I find no evidence that *Wine Spectator* allows advertising directly to influence which wines receive awards.

V. Conclusion

With the notable exceptions of publications like *Consumer Reports* and *Wine Advocate*, which are funded entirely by subscriptions, the majority of the publications that provide

consumers with product reviews accept advertising from the very firms whose products they review. In this paper, I test whether *Wine Spectator* rewards its advertisers with higher wine ratings or more awards. Put differently, I test whether—within this market—the returns to a reputation for producing unbiased reviews outweigh the returns to bundling favorable reviews with advertising.

Overall, I find little consistent evidence of bias. At worst, the tests for biased ratings suggest that *Wine Spectator* rates wines from advertisers almost one point higher than wines from non-advertisers. However, selective retastings can explain at most half of this bias and then only within the set of U.S. wines rated by both *Wine Spectator* and *Wine Advocate*. Given *Wine Spectator's* claim that it rates wines blind, the remaining difference in ratings may simply reflect consistent differences in how the two publications rate quality, which leads to predictable differences in advertising. This interpretation is consistent with the fact that tests for biased awards provide no additional evidence of bias. Therefore, despite the fact that *Wine Spectator* is dependent on advertising revenue, the long-run value of producing credible reviews appears to minimize bias.

The fact that *Wine Spectator* reviews are largely without bias may have implications for publications in other markets as well. After all, there were several reasons to expect that wine reviews might be biased. First, because the costs of biased wine reviews are low relative to the costs of biased reviews for other types of products, the typical consumers might be relatively tolerate of biased wine reviews. Second, the fact that a subset of consumers are willing to fund *Wine Advocate* entirely through subscriptions is consistent with these consumers believing that advertising influences the wine reviews published in other publications, and with other consumers not valuing unbiased reviews enough to pay the markedly higher price of *Wine Advocate* measured on a per review basis. Third, there is a subjective component to rating wine that does not exist, for example, when *Morningstar* uses mutual fund returns to rank mutual funds. Therefore, when the costs of biased reviews are relatively high or when it would be relatively easy for consumers to determine that a publication has biased its reviews, the results in this paper suggest that reputational considerations may keep publications honest. (Of course, not all product reviews lend themselves to blind comparisons.)

Alternatively, the lack of biased wine reviews might reflect the fact that the returns to receiving biased wine ratings are relatively low. For example, the availability of *Wine Advocate* ratings in wine shops may reduce the weight consumers place on *Wine Spectator* ratings, thereby lowering a winery's return to a biased *Wine Spectator* rating. Regardless, this discussion highlights an important limitation of my empirical results. When confronted with large financial returns to biasing their opinions, experts like accounting firms, securities analysts, ratings agencies, and members of the financial media may be little constrained by reputational considerations.

References

Ackerberg, D.A. (2001). Empirically Distinguishing Informative and Prestige Effects of Advertising. *RAND Journal of Economics*, 32, 316–333.

Amemiya, T. (1985). *Advanced Econometrics*. Cambridge, Massachusetts: Harvard University Press.

Ashenfelter, O. (1989). How Auctions Work for Wine and Art. *Journal of Economic Perspectives*, 3, 23–36.

Ashenfelter, O., Ashmore, D. and Lalonde, R. (1995). Bordeaux Wine Vintage Quality and the Weather. *Chance*, 8, 7–14.

Ashenfelter, O. (2008). Predicting the Quality and Prices of Bordeaux Wine. *Economic Journal*, 118, F174–184.

Ashenfelter, O. and Storchmann, K. (2009). Using Hedonic Models of Solar Radiation and Weather to Assess the Economic Effect of Climate Change: The Case of Mosel Valley Vineyards. *Review of Economics and Statistics* (forthcoming).

Byron, R.P. and Ashenfelter, O. (1995). Predicting the Quality of an Unborn Grange. *Economic Record*, 71, 40–53.

Darby, M.R. and Karni, E. (1973). Free Competition and the Optimal Amount of Fraud. *Journal of Law and Economics*, 16, 67–88.

Del Guercio, D. and Tkac, P. (2008). Star Power: The Effect of Morningstar Ratings on Mutual Fund Flow. *Journal of Financial and Quantitative Analysis*, 43, 907–936.

DellaVigna, S. and Gentzkow, M. (2009). Persuasion: Empirical Evidence. *Annual Review of Economics* (forthcoming).

Nelson, F. and Olson, L. (1978). Specification and Estimation of a Simultaneous-Equation Model with Limited Dependent Variables. *International Economic Review*, 19, 695–709.

Nelson, P. (1970). Information and Consumer Behavior. *Journal of Political Economy*, 78, 311–329.

Nelson, P. (1974). Advertising as Information. *Journal of Political Economy*, 82, 729–754.

Newey, W. (1987). Efficient Estimation of Limited Dependent Variable Models with Endogenous Explanatory Variables. *Journal of Econometrics*, 36, 231–250.

Reuter, J. and Zitzewitz, E. (2006). Do Ads Influence Editors? Advertising and Bias in the Financial Media. *Quarterly Journal of Economics*, 121, 197–227.

Reinstein, D.A. and Snyder, C.M. (2005). The Influence of Expert Reviews on Consumer Demand for Experience Goods: A Case Study of Movie Critics. *Journal of Industrial Economics*, 53, 27–51.

Stigler, G.J. (1961). The Economics of Information. *Journal of Political Economy*, 69, 213–225.