1-2016

Wintertime for Deceptive Advertising?

Jonathan Zinman
Dartmouth College

Eric Zitzewitz
Dartmouth College

American Economic Journal: Applied Economics 2016, 8(1): 177–192
http://dx.doi.org/10.1257/app.20130346

Wintertime for Deceptive Advertising?[†]

By Jonathan Zinman and Eric Zitzewitz*

Casual empiricism suggests that deceptive advertising about product quality is prevalent, and several classes of theories explore its causes and consequences. We provide unusually sharp empirical evidence on its extent, mechanics, and dynamics. Ski resorts self-report substantially more natural snowfall than comparable government sources. The difference is more pronounced on weekends, despite third-party evidence that snowfall is uniform throughout the week—as one would expect given plausibly greater returns to exaggeration on weekends. Exaggeration is greater for resorts that plausibly reap greater benefits from it: those with expert terrain and those not offering money back guarantees. (*JEL* D83, L15, L83, M37, Z31)

Casual empiricism suggests that deceptive advertising about product quality is prevalent, and several classes of theories explore its causes and consequences.[1] Yet there is little sharp empirical evidence that speaks to such theories. This gap is due in part to formidable measurement challenges; in most settings, measuring deceptive advertising requires detailed, high-frequency information (on ads, product quality, and inventories) that is difficult to observe.

We test for deceptive advertising by examining a critical component of product quality at ski resorts: *new*, *natural* (or "fresh") snowfall in the past 24 hours. Ski resorts issue "snow reports" on their websites roughly once a day. These reports are also collected by aggregators and then rebroadcast over the Internet and via print and broadcast media. A skier wishing to ski on new, natural snow can use these snow reports to help decide whether and where to ski on a particular day. In principle, snow reports provide skiers with location-specific information on fresh snowfall that is not necessarily captured by third-party weather websites.

*Zinman: Department of Economics, Dartmouth College, HB 6106, Hanover, NH 03755 (e-mail: jzinman@dartmouth.edu); Zitzewitz: Department of Economics, Dartmouth College, HB 6106, Hanover, NH 03755 (e-mail: eric.zitzewitz@dartmouth.edu). Thanks to Noah Glick, Meagan Herzon, Percy Lee, Benjamin Lo, and Boris Vabson for outstanding research assistance, and to James Feyer, Thomas Hubbard, Christine Jolls, Chris Snyder, Michael Spence, Kip Viscusi, two anonymous referees, and seminar audiences at Dartmouth College, the Utah Business Economics Conference, and the National Bureau of Economic Research (NBER) Law and Economics Summer Institute meeting for comments. The views expressed in this paper are not necessarily those of our employer, which owns a small resort in our sample.

† Go to http://dx.doi.org/10.1257/app.20130346 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

1 Lazear (1995) models firms taking advantage of high consumer search costs by "baiting" consumers with a high quality good that is not actually in-stock, and then "switching" consumers to a lower quality good that is in-stock. In signal-jamming models (e.g., Holmstrom 1999), agents engage in costly effort to upwardly bias signals of their quality, but rational recipients of these signals anticipate these efforts and no information is lost. Other models focus on deception more generally, with motivating examples from advertising; e.g., Ettinger and Jehiel (2010). See also footnote 10.

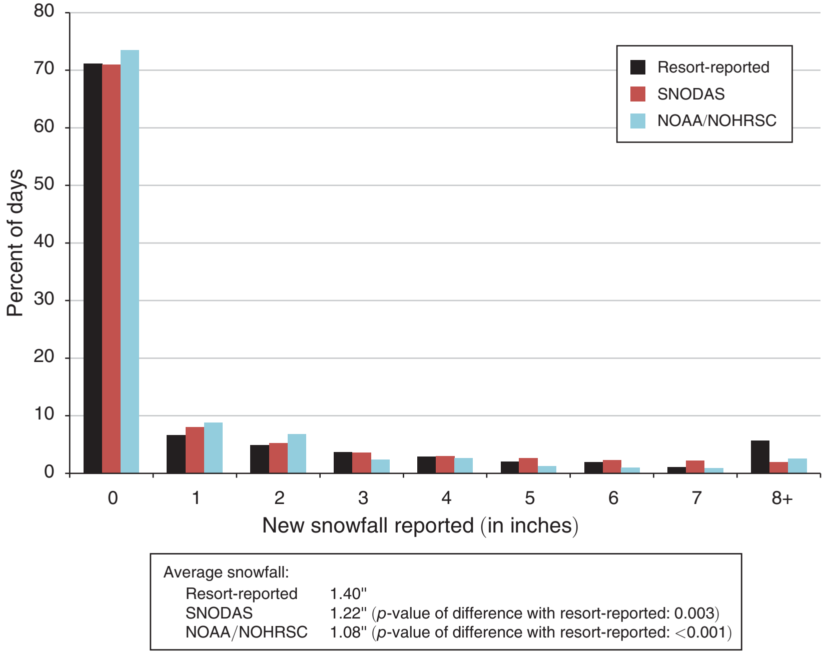

Figure 1. Distribution of Snow Reports, Resort-Reported versus Government

We start by estimating that ski resorts report 12–23 percent more snow than government sources during our sample period (covering the 2004–2005 through 2007–2008 ski seasons). This holds even when the government comparison is based on a National Weather Service model (SNODAS) that plausibly captures the locational advantages of resort microclimates (Figure 1).

Allowing for the possibility that the Snow Data Assimilation System (SNODAS) does not fully capture the locational benefit of resorts, we also test for exaggeration using *within*-resort variation in reporting behavior over the course of the week, based on the premise that the returns to any deceptive advertising will be greater on weekends. Resorts only benefit from exaggerating snow reports when skiers base their purchase decisions on them. The cost of exaggeration is angering or losing credibility with skiers, including those who have already precommitted (e.g., as part of a multiday vacation) but use the snow report to help plan their day. Several factors suggest that resort snow reporting behavior could differ on weekends versus weekdays. Customer traffic tends to be much heavier on weekends (as evidenced by resorts' extensive use of congestion pricing),[2] so resorts may tend to exaggerate more on weekends if the benefits of exaggeration scale with demand while the cost—which we suspect is primarily reputational—scales less.[3] Likewise, resorts

[2] See, e.g., http://www.tripadvisor.com/ShowTopic-g33701-i1438-k3129047-When_Is_The_Best_Time_To_Ski-Winter_Park_Grand_County_Colorado.html on avoiding lift lines by skiing during the week, and Fry (2008) on peak versus off-peak pricing.

[3] There may also be factors pushing toward greater exaggeration on weekdays; e.g., if there are more expert and hence fresh snow-elastic skiers during the week (a "ski bum" effect, if you will). We lack the requisite demand-side

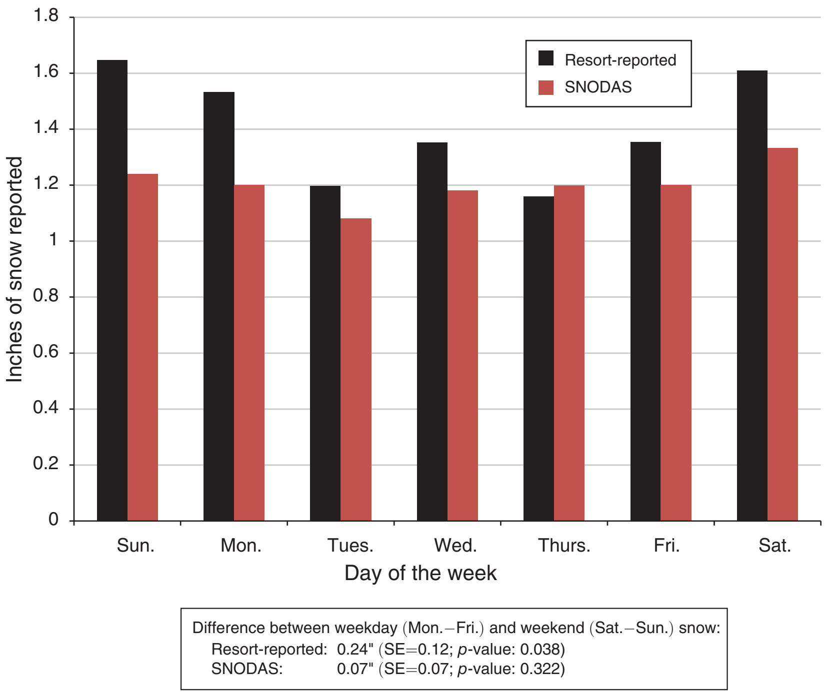

FIGURE 2. SNOWFALL BY DAY OF WEEK, RESORT-REPORTED VERSUS GOVERNMENT

may exaggerate more on weekends if there is a larger (potential) share of precommitted customers on weekdays than on weekends, when many skiers are less constrained by work schedules and hence more likely to condition resort choice on snow conditions.[4]

We estimate that resorts report about 22 percent more new natural snow on Saturday and Sunday mornings (difference of 0.24″, p-value = 0.04). This pattern holds despite the fact that both climatology research[5] and government measures (Figure 2) suggest that snowfall is uniform throughout the week. The resort-reported "weekend effect" turns out to be substantial in absolute as well as percentage terms; an extra quarter-of-an-inch might not sound like much, but we find evidence that the weekend exaggeration is concentrated on the intensive margin, that it only snows on about 30 percent of the days in our sample, and that many resorts report accurately

data to formulate a sharp hypothesis about whether exaggeration is more likely on weekends or weekdays, and hence conduct nonparametric and two-sided tests for whether and how much snow reporting varies over the week.

[4] Conversely, skiing decisions are probably relatively inelastic with respect to snow during the week because work schedules are the dominant decision factor in how and where most prospective skiers spend their weekday daytimes.

[5] There is a literature exploring the possibility that pollution drives intraweek variation in weather, but the effects sizes estimated in that literature are quite small (relative to the 12–23 percent difference in resort-reported snow we find), of mixed sign, and concentrated in nonskiing areas. Cerveney and Balling (1998) find higher CO and O_3 levels and higher precipitation on Fridays and Saturdays. Effects are primarily in areas downwind of the US Eastern seacoast. Forster and Solomon (2003) find that nighttime low temperatures are 0.2 to 0.4 degrees Celsius higher on weekends in the middle of the United States but are 0.1 to 0.2 degrees lower in the Southeast and Southwest (weekend effects in the Northeast and West, where most of our resorts are located, are smaller).

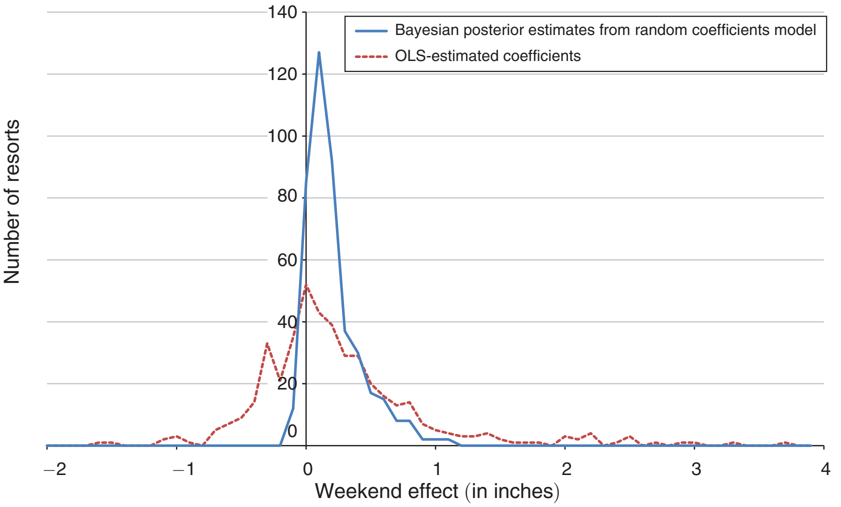

Figure 3. Distribution of Resort-Level Estimates of Weekend Effects

(Figure 3). Overall then, our results suggest that, when resorts exaggerate, they do so by an inch or more.[6] The weekend effect persists even if we take the conservative step of controlling for government-reported snow.

Having found some evidence that deceptive advertising varies *within* resorts over time along with payoffs, we next explore whether deceptive advertising varies *across* resorts with plausibly different payoffs. We find evidence that both weekday and weekend exaggeration in resort reporting are larger for resorts with more expert terrain and for those that do not offer a money back guarantee.[7] This is consistent with expert skiers valuing fresh snow more highly, and with guarantees and deception being substitutes (rather than, e.g., guarantees giving skiers a false sense of confidence in resort reports). We also find some evidence that deceptive advertising increases with competition and market size (population within 150 miles), although these results are not robust to controlling for alternative government measures of snowfall. We do not find evidence that weekend effects vary by resort ownership type (government, privately held firm, or publicly-traded), although these null effects are imprecisely estimated.

In all, our results suggest that deceptive advertising about product quality responds sharply to incentives, both within-resort over time—as evidenced by the weekend effect—and across resorts with different characteristics.

[6] In addition to reporting new natural snow, resorts also report other aspects of snow quantity and quality, such as base depth, number of trails open, and surface conditions (e.g., powder, packed powder). These other aspects can be influenced by manmade snow, and hence we do not examine them.

[7] Guarantees, when offered, typically offer the skier a 100 percent refund on that day's lift ticket so long as the skier does not take more than one run.

Our work is related to literature on how customer feedback on product/service quality affects firm behavior and equilibrium.[8] It also complements prior work showing that third-party quality disclosure can change firm behavior.[9] Our finding of significant product quality exaggeration builds on Jin and Kato (2006), who audit claims about the quality of baseball cards being auctioned on eBay and find evidence of exaggeration.[10] Jin and Kato's analysis focuses on the effect of deceptive claims on demand and auction prices, whereas our focus is more on the supply of deception.[11]

Although our findings provide unusually sharp evidence on the nature and dynamics of deceptive advertising, our setup can only identify a subset of the behaviors of interest for modeling and policy analysis. We lack any high-frequency data on prices or quantities sold, and hence cannot measure the demand responses to snow reports.[12] This prevents us from sharply testing across the models discussed at the outset or measuring the welfare implications of (changes in) advertising practices. Moreover, as with any study of a single industry, the external validity of our findings to other markets is uncertain.

Our findings nevertheless have potentially broad applicability, as the market for skiers does not seem uniquely suited to deceptive advertising. There are many other markets where search and switching costs loom large.[13] And many of the other conditions that contribute to deception in theory do *not* seem to prevail in our setting. Ski area customers get immediate and visceral feedback on the accuracy of snow reports. The potential for repeat play and learning is high.[14] And the entry and exit rates of ski areas are low: there are few, if any, "fly-by-night" players with incentives to commit outright fraud. So we speculate that there are many other markets where conditions are ripe for deceptive advertising that varies sharply with advertiser incentives.

The paper proceeds as follows. Section I details our data on snowfall and on resort characteristics. Section II details our identification strategy and results. Section III concludes.

[8] This literature has focused to a great extent on eBay; see, e.g., Cabral and Hortaşçu (2010) and citations therein. See also Hubbard (2002) on auto repair and Luca (2011) on restaurants.

[9] See, e.g., Sauer and Leffler (1990), Dranove and Jin (2010), and Luca and Smith (2015). Beales, Craswell, and Salop (1981) provide an overview of legal limitations on deceptive advertising. Jolls and Sunstein (2006) discuss behavioral motives for government regulation of advertising.

[10] See also Ellison and Ellison (2009) for evidence of firm practices that "frustrate consumer search or make it less damaging to firms" in the Pricewatch shopping engine for computer parts. Luca and Zervas (2013) and Mayzlin, Dover, and Chevalier (2014) identify and analyze review fraud in restaurants and hotels; in their settings, unlike ours, the sender of the (deceptive) message about product quality is not readily identifiable.

[11] Another related literature considers whether, how, and why entities "game" mandated reporting re: product quality; see, e.g Forbes, Lederman, and Tombe (2015) on airline reporting of on-time performance, and Chen et al. (2012) on air pollution data in China.

[12] We conducted this study without industry cooperation, and hence our data on resorts is limited to what we could gather from other sources. Nearly all resorts are privately held so publicly available financial information on them is scarce, particularly at the daily frequency.

[13] In our setting, on a one-shot basis, driving times are substantial even between "neighboring" resorts. On a longer-term basis, some consumers may find it costly to coordinate with peers on alternative destinations or to learn how to navigate the terrain and ancillary services (parking, rentals, dining, lodging) of a new mountain.

[14] The immediate feedback stands in contrast to the examples (e.g., tobacco use, investment advice) that motivate Glaeser and Ujelyi (2010), and Kartik, Ottaviani, and Squintani (2007). The visceral feedback (and high stakes) contrasts with the "low involvement situations" (e.g.,. voting, cheap products) that can make consumers susceptible to persuasion in Mullainathan, Schwartzstein, and Shleifer (2008).

I. Snowfall and Resort Data

Our data for measuring product quality and its reporting consists of resort-provided snow reports, government snow data, and resort characteristics.

Resort-reported snow is measured by the resorts themselves, and reported, typically daily, on individual resort websites. Our discussions with industry participants suggest that snow measurement practices vary across resorts: some take multiple measurements up-and-down the mountain, while others measure from a single spot. One commonality seems to be that there is a division of labor: there are one or more employees who take the measurements and then call/radio in to the marketing department. Marketing is then responsible for reporting snow conditions on the resort's website. These practices suggest that resorts have several potential production functions for exaggerating fresh snowfall, and for doing so selectively on weekends. Measurement-takers might be more likely to fluff or round on the weekend (e.g., reporting 3 inches instead of 2, or 2 inches instead of 1.4). Marketers might be more likely to fluff or round numbers, reported to them by measurement-takers, on the weekend. Marketers could instead/also take a higher number from among the many provided by measurement-takers from different parts of the resort; e.g., marketers could report a mid-mountain number during the week, and a summit number on the weekend. We test the hypothesis that resorts report more snow on weekends, without being able to identify the exact production function of any exaggeration.

We collect resort-provided reports, for the 2004–2005 through 2007–2008 ski seasons, from the websites of two popular aggregators: *Skireport.com* and *Onthesnow.com*.[15] These websites do not supply archives of ski reports, and thus we assembled the data from different sources. Our primary source is two private Internet archives. Since these archives had better coverage of *Onthesnow.com*, we used archived reports from this website. In the data collected from Internet archives, we are limited to collecting data for days on which the relevant web page was archived. We collect snow reports from archived pages that summarize all reports from a given state or province, so an entire state's data is either archived or not on a given day. The frequency of data collection in these archives increases over time. In the 2004–2005 and 2005–2006 seasons, snow reports are available for only about 10 percent of resort-days between December and March. This ratio rises to 30 percent in the 2006–2007 season and 65 percent in the 2007–2008 season.[16]

The archiving process is Internet-wide, so it seems reasonable to assume that the archiving of data for a resort should be exogenous to actual or reported snow. We test this assumption in three ways. First, we test whether reports were more likely to be archived on certain days of the week, and find that weekends account for almost exactly two sevenths of our resort reports (28.4 percent, p-value of difference from $2/7 = 0.945$). Second, we simply examine the timing of the reports, finding that in one of our archives it increased from once every ten days, to once

[15] Since we circulated the first draft of this paper in mid-2009, both *Skireport.com* and Mountain News Corporation, the owner of *Onthesnow.com*, have been purchased by Vail Resorts. *Skireport.com* now redirects to *Onthesnow.com*.

[16] Starting February 15, 2008 we also collected snow reports once per day from *Skireport.com*.

every five days, to once every three days, to essentially every day. Archiving frequencies were higher for states with more resorts (e.g., Colorado versus Missouri), suggesting that the archiving of a page containing a state's snow reports responded to that webpage's popularity, but the regular sampling frequencies suggested that this response was not happening at high enough frequency to contribute to a weekend effect. Third, we test whether the availability of a report is correlated with an interaction of government-reported snow (which we can measure on days when resort-reported snow is missing from the archives) and a weekend indicator variable, and find no evidence that it is (results discussed below, and reported in online Appendix Table A1).

Our government data on snowfall comes from two sources: actual reported snowfall from nearby government weather stations, and estimated snowfall from the Snow Data Assimilation System (SNODAS), a US National Weather Service model that provides estimated snowfall from satellite, ground station, and airborne weather data collection.[17] SNODAS data are available for any point in the continental United States on a 30-arc-second grid.[18] We take the largest of the 25 SNODAS estimates from the 5×5 grid surrounding the main resort mountain as the estimate of actual snowfall that we match to the resort snow report.[19]

For the government weather stations, we match each resort with to up to 20 National Oceanographic and Atmospheric Administration (NOAA) or National Operational Hydrologic Remote Sensing Center (NOHRSC) weather stations within 100 miles horizontally and at elevations within 1,000 feet of the resort summit.[20] We match each resort snow report to mean reported snow from the surrounding stations that meet these criteria.

In matching the resort and government snow data, we match time periods as closely as possible. Resorts can issue and update snow reports on aggregator websites at any time, but they usually issue a report early in the morning local time. This report is timed to capture as much overnight snowfall as possible while still being available in time to affect that day's skier purchasing decisions. Saturday's snow report issued at 7 AM local time would therefore reflect snowfall from 7 AM Friday to 7 AM Saturday, and so we attempt to match the Saturday resort report with SNODAS and weather station data from this time period. NOHRSC reports typically cover a 24-hour period beginning at 7 AM local time, so this matches the timing of resort reports well. NOAA stations aggregate their data into 24-hour periods beginning at midnight Coordinated Universal Time (UTC), which corresponds to 7 PM Eastern Standard Time (EST) and 4 PM Pacific Standard Time (PST) in winter months. Since NOHRSC reports provide a better match with the timing of the resort reports,

[17] SNODAS data are described and available at http://nsidc.org/data/g02158.html.

[18] Thirty arc seconds are roughly 930 meters north-south and 660 meters east-west (at 45 degrees latitude).

[19] We collect data on the latitude and longitude of resort mountains primarily from the US Geological Survey, and supplement this data with hand-collected information from Google Maps.

[20] NOAA station data are described at http://www.ncdc.noaa.gov/oa/climate/ghcn-daily/. NOHRSC station data is described at http://www.nohrsc.noaa.gov/nsa/reports.html. We match weather stations using a loss function of the distance in miles plus 0.1 times the difference between summit and station elevation in feet. The average matched station is 26 miles away and 160 feet below the summit for Eastern resorts and 52 miles away and 280 feet below the summit for Western resorts. Twenty-eight out of 437 resorts do not have matching weather stations due to the elevation restriction (19 of these are in Western Canada).

we match with NOAA reports only if matched NOHRSC stations are not available. SNODAS aggregates its data into 24-hour periods that begin at 6 AM (UTC), or 1 AM EST and 10 PM PST. Our analysis accounts for this asynchronicity.

Table 1 provides statistics on snowfall, resort characteristics, and the distribution of our resort and snowfall observations across region and time period. We include resorts in the United States and Canada in our sample. Forty-two percent of resorts have any expert terrain (black or double-black diamond); 59 percent offer a money-back guarantee; and 90 percent are privately owned. The mean number of competitors (resorts within 50 miles as the crow flies) is 5.9, with a standard deviation of 4.8. SNODAS forecasts are not available in Canada, but are available for essentially all US resorts. As mentioned above, the frequency of resort snow reports increases later in the sample, from about 12–13 reports per resort × year in 2004–2005 and 2005–2006, to 39 in 2006–2007, to 70–75 in 2007–2008.

II. Results

A. *Resorts Report More Snow than Government Sources*, *on Average*

Figure 1 shows that resorts report more snow than government sources on average: 1.40 inches per resort-day versus 1.22 (SNODAS) or 1.08 (NOAA/NOHRSC). The *p*-values on these differences are 0.003 and <0.001.[21] The difference is driven by resorts reporting substantially more days with eight or more inches of snow. To be fair, even SNODAS may not fully capture the benefit of resorts locating themselves to take advantage of additional snowfall generated by very local microclimates,[22] and so we proceed with tests for deceptive advertising that do not rely on resort versus government snowfall comparisons.

B. *Identifying a Weekend Effect: Resorts Report More Snow on Weekends*

Figure 2 illustrates that resorts report more snow on weekends than weekdays, with the lowest levels in the middle of the week and the highest on Saturday, Sunday, and Monday.[23] Meanwhile, government reports are relatively flat throughout the week (the *p*-value on the weekend effect indicator is 0.32).[24]

More formally, our starting point for testing for weekend effects is the OLS specification:

$$s_{rt} = \beta \times w_t + \mathbf{a}_w + \mathbf{n}_r + e_{rt}, \tag{1}$$

[21] The *p*-values given, like the standard errors in the regressions that follow, allow for clustering of error terms within days and resorts.

[22] Although the SNODAS model estimates snowfall for precise locations, it does so partly by interpolating between government weather stations. Thus, strategic ski resort location decisions might exploit mountain-by-mountain variation in snowfall that SNODAS does not fully capture.

[23] Note that both Martin Luther King, Jr. Day and Presidents' Day fall on Mondays and are part of popular long weekends for skiing. Online Appendix Table A2 presents estimates of the resort-reporting weekend effect for various subperiods (by year, by month, and by holiday versus nonholiday). We do not have sufficient power to draw any firm inferences.

[24] The daily averages in Figure 2 are for the sample used in the regressions in Table 2 and the differences from the regression in Table 2, column 2 (and an analogous regression with SNODAS on the left hand side).

Table 1—Resort Characteristics and Sample Size

Panel A. Resort characteristics	All resorts (437) Mean	SD	
Terrain type (percent)			
Beginner	27	11	
Intermediate	41	12	
Advanced	25	13	
Expert	7	11	
Percent with any expert terrain	42		
Base elevation	3,132	2,946	
Summit elevation	4,373	3,580	
Vertical drop	1,242	1,004	
Lifts	7.8	5.1	
Runs	41	35	
Acres	671	1,114	
ln (population within 150-mile radius)	15.5	1.2	
Competing resorts within 50 miles	5.9	4.8	
Percent offering money-back guarantee	59		
Type of ownership (percent)			
Publicly traded	3.2		
Private	89.9		
Government	6.9		
Panel B. Sample size	Reports	Unique resorts	Unique days
Resort-generated reports	56,402	437	752
With matched NOAA-NOHRSC weather station data	47,741	417	709
With matched SNODAS data (US only)	39,920	320	707
By ski season			
2004–2005	4,054	354	171
2005–2006	4,807	363	176
2006–2007	15,376	393	201
2007–2008	32,165	429	204
By region			
Eastern United States	24,390	232	594
Western United States	24,825	135	693
Eastern Canada	3,132	41	311
Western Canada	4,055	29	477
Panel C. Average snowfall	Mean	SD	
Resort-reported	1.40	3.42	
SNODAS	1.22	1.78	
NOAA/NOHRSC	1.08	2.32	
Observations for which all three are available	35,965		

Note: States and provinces that are entirely east of the Continental Divide are considered eastern.

where s_{rt} is natural new (or "fresh") snowfall reported by resort r on day t, w_t is an indicator variable for whether t is a weekend day, $\mathbf{a}_w$ is a vector of fixed effects for calendar weeks (Wednesday–Tuesday), $\mathbf{n}_r$ is a vector of resort fixed effects, and e_{rt} is an error term. The fixed effects control for any bias arising from the proportion of snow reports on weekends varying between more and less snowy weeks of the year (e.g., if resorts were open only on weekends at the beginning and end of the season) or between more and less snowy resorts.[25]

[25] Point estimates actually change very little when these fixed effects are dropped, so these potential omitted variable biases do not appear important in practice. But the fixed effects do improve the efficiency of estimation

Since actual and reported snowfall may be correlated across resorts on the same day, we allow for clustering within day when calculating our standard errors. Since snow reports may be serially correlated, we also allow for clustering within resorts, using the two-dimensional clustering procedure in Petersen (2009). Allowing for clustering within days does meaningfully affect standard errors, while clustering within resort has essentially no effect. Clustering at other levels, or allowing for spatial correlation in the error term, tends to produce slightly smaller standard errors than our preferred method of clustering on resort and day (online Appendix Table A3).

The bottom of Figure 2 reports the key result of equation 1's test for a weekend effect in resort reports: a statistically significant effect of 0.24 inches (p-value $= 0.04$). There is no comparable effect in SNODAS, although the weekend point estimate is positive at 0.07 inches (p-value $= 0.32$).

Table 2 shows some robustness checks for the resort-reported weekend effect (see also online Appendix Tables A2–A4). Column 1 shows we get a similar result if we estimate equation (1) on the full sample of resort-days instead of just those with SNODAS data. Column 2 is our main specification and reproduces the result reported at the bottom of Figure 2. Columns 3–5 add controls for various functional forms of SNODAS-measured snowfall: column 3 includes same-day SNODAS, column 4 weights SNODAS observations across days to account for timing differences between resort and government reports, and column 5 flexibly controls for leads and lags (longer leads and lags do not have statistically significant coefficients). Regardless of specification, adding SNODAS controls reduces both the standard error and the point estimate on the weekend effect but leaves it economically and statistically significant (online Appendix Table A4 shows a similar pattern if we use weather station data instead of SNODAS).

Online Appendix Table A1 examines whether our results are affected by either selective reporting by resorts or by the selection of reports captured in our archives. Selective reporting could be due to (possibly also deceptive) practices by resorts or aggregators that are subtly different than reporting more snow than has actually fallen. For example, resorts with no new snow might prefer to leave up a stale report rather than report no new snow, and the optimality of this strategy might vary over the week (columns 4–6). Similarly, we might be concerned that archiving was related to interest in a webpage, and thus might be related to snow in a way that varied over the week (columns 1–3). In practice, we find some evidence that reporting and archiving frequencies are very slightly higher when there is more actual snow (columns 2, 5, and 8). But these effect sizes are tiny (compared to the dependent variable mean in the last row), and they disappear once we include our usual set of fixed effects for weeks and resorts (columns 3, 6, and 9).[26] More importantly, there

by absorbing variation that would otherwise be captured by the error term. Results are also qualitatively similar if we use Tobit instead of OLS. We also estimate specifications with resort*week fixed effects. This reduces degrees of freedom by about a third and increases the standard errors accordingly. The weekend effect point estimate is not statistically (or economically) significantly different from the specifications in Table 2 and the additional fixed effects are not collectively significant.

[26] In particular, it is the inclusion of resort fixed effects that causes the relationship between snow and both reporting and archiving frequencies to disappear, suggesting that this relationship reflects snowier resorts reporting more frequently and snowier states being archived more frequently.

Table 2—Weekend Effect Regressions, with and without Controlling for Actual Snowfall

Dependent variable: Inches of new natural snowfall reported by resort					
Observations include	All observations	w/SNODAS			
	(1)	(2)	(3)	(4)	(5)
Weekend (Sat. and Sun.)	0.233**	0.242**	0.174*	0.175*	0.183**
	(0.099)	(0.117)	(0.101)	(0.091)	(0.089)
SNODAS $(t+1)$					0.359***
					(0.025)
SNODAS (t)			0.716***		0.621***
			(0.042)		(0.029)
SNODAS $(t-1)$					0.104***
					(0.020)
Weighted SNODAS				0.967***	
				(0.054)	
Observations	56,402	39,920	39,920	39,920	39,920
Unique days	752	707	707	707	707
R^2	0.123	0.141	0.309	0.341	0.351

Notes: Each specification is estimated by OLS with fixed effects for weeks (Wed.–Tues.) and resort. Weighted SNODAS averages the SNODAS observations for days t and $t+1$, weighting by the number of hours in the 7AM-to-7AM local time window that overlap with the SNODAS observation window in question (weights are 0.75 on t and 0.25 on $t+1$ for the Eastern time zone; and 0.625 and 0.375, respectively, for the Pacific time zone). Columns 2–5 are restricted to observations with SNODAS data for days $t-1$ to $t+1$. Standard errors allow for clustering within both day and resort.

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

is no evidence that selectivity differs on weekends (this key result is shown in the "Gov. snow × Weekend" row).[27]

C. *Margins and Functional Form of Exaggeration*

Table 3 explores resort reporting behaviors vis-à-vis the functional form of exaggeration. Here we take equation (1) and replace the single weekend indicator with several interactions between the weekend indicator and indicators of different levels of SNODAS snowfall (the regressions also include the uninteracted SNODAS indicators, but we do not report these in the table). We do this for the extensive versus intensive margins of resort reporting (columns 1 and 3), and for our usual measure that combines both margins (column 2).

As noted previously, Figure 1 suggests that resorts may be exaggerating conditional on receiving a large amount of snow, so the bottom row of results in Table 3 shows the test for a weekend effect conditional on our SNODAS estimate

[27] We were also concerned that resorts may issue updated snow reports on days when snow is falling, and that this updating may be more common on weekends. In practice, however, updated snow reports are quite rare. Of the 1504 times our third-party archives captured the same resort twice on one day, only 62 times (4.5 percent) had the resort report been updated, and these instances were not disproportionately on weekends or weekdays. If same-day decisions about where to ski are largely made early in the morning, there may be limited incentive to update midday.

Table 3—Extensive versus Intensive Margin

Dependent variable:	= 1 if resort reported snow > 0	Resort-reported snow (inches)	
Sample:	Full (1)	Full (2)	Resort snow > 0 (3)
Indicator variables for levels of SNODAS snow			
SNODAS = 0.01 to 0.49″	0.104***	0.447***	0.290
	(0.009)	(0.061)	(0.271)
SNODAS = 0.50 to 0.99″	0.265***	1.006***	0.737**
	(0.015)	(0.091)	(0.290)
SNODAS = 1.00 to 5.99″	0.496***	2.793***	2.543***
	(0.017)	(0.151)	(0.297)
SNODAS = 6.00+	0.698***	6.363***	5.298***
	(0.024)	(0.431)	(0.414)
Interactions of indicator variables with weekend effect			
Weekend × (SNODAS = 0)	0.004	0.044	0.208
	(0.007)	(0.058)	(0.469)
Weekend × (SNODAS = 0.01 to 0.49″)	0.036***	0.148**	0.069
	(0.013)	(0.071)	(0.242)
Weekend × (SNODAS = 0.50 to 0.99″)	0.028	0.142	0.181
	(0.023)	(0.114)	(0.227)
Weekend × (SNODAS = 1.00 to 5.99″)	0.019	0.395	0.330
	(0.020)	(0.246)	(0.268)
Weekend × (SNODAS = 6.00+)	−0.002	0.318	0.098
	(0.034)	(0.745)	(0.676)
Observations	39,920	39,920	11,523

Notes: This table reports the results of regressions that are identical to Table 2, column 2, except that the single weekend indicator is replaced by interactions of that variable with indicator variables for different levels of weighted SNODAS snowfall, and the regressions also include the uninteracted indicator variables. Regressions also include the resort and week (Wed.–Tue.) fixed effects as in Table 2. Standard errors allow for clustering on day and resort.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

being six inches or more. We find only weak support for that hypothesis (see also online Appendix Table A5, where we use weather station instead of SNODAS data).

Table 3 shows stronger evidence that resorts are more likely to report a positive amount of snow, when government reports are at levels of a trace to $< 1/2$ of an inch, on weekends than on weekdays (columns 1 and 2). This is likely due to resorts being more likely to round up fractional amounts on weekends, since we never observe a resort reporting a noninteger amount of snow in the (0, 1) interval. Online Appendix Table A5 shows a similar but statistically stronger pattern using weather station data instead of SNODAS as a benchmark.

All told, these results suggest the weekend effect may be due to two distinct types of exaggeration. First, resorts seem more likely to round a fractional inch of snow up to a single inch on weekends. Second, resorts may inflate amounts when a large amount of snow falls.

D. *The Weekend Effect and Resort Characteristics: Which Resorts Exaggerate?*

Next we estimate cross-resort differences in weekend effects. Figure 3 plots the distribution of resort-level weekend effect estimates, with the resort count on the y-axis. The solid line (with the larger mass near zero) comes from a Bayesian posterior mixed (random coefficients) model; the dashed line comes from a noisier model that interacts the weekend indicator with resort fixed effects.[28] Both models suggest that the modal weekend effect is close to zero, and that a substantial number of resorts (perhaps 10 to 15 percent) exaggerate by 0.5 inches or more on weekends.

Table 4 tests whether the weekend effect is larger for resorts with characteristics that proxy for payoffs to exaggeration by adding those proxies, and interactions between those proxies and the weekend dummy, to the specification we report in Table 2 column 4. We find that resorts with expert terrain report about 0.14 inches more on weekends than resorts without expert terrain (columns 2–4).[29] We also find evidence of larger weekend effects for resorts that do not offer money-back guarantees (i.e., the interactions between weekend and no guarantee are positive and significant).[30] The finding that exaggeration and guarantees are not complementary suggest that guarantees increase the expected cost of exaggeration (by inducing consumers to exercise guarantees) more than they increase the expected benefit of exaggeration (by, e.g., giving consumers a false sense of security). We find the same pattern of results on expert terrain and guarantees if we control for weather station snow instead of SNODAS (online Appendix Table 6). Table 4 also shows some evidence of greater weekend exaggeration by resorts in more densely populated markets and resorts facing more competition,[31] although these inferences are not robust to controlling for weather station snow instead of SNODAS (online Appendix Table 6). We find no evidence that region (east or west of the Continental Divide) or ownership (public, government, or private) mediate weekend effects, although these estimates are noisy nulls.

The results on expert terrain in particular raise further questions about the functional form of the weekend effect, since expert resorts have a substantially higher *level* of snowfall. One might think to do this exercise using resort/government snow as an outcome variable in a regression, but this would be problematic given the many resort-days with trace amounts of snow in the government data because resorts report integer amounts of snow in 99.8 percent of cases, and in 100 percent

[28] We estimate the mixed model using the Generalized Linear Latent and Mixed Model (GLLAMM) procedure described in Rabe-Hesketh, Skrondal, and Pickles (2004). Both the GLLAMM model and the fixed effects model reject the null hypothesis of a constant weekend effect across resorts (with a p-value less than 0.0001 in both cases).

[29] We also estimate specifications that interact the share of terrain that is intermediate, advanced, and expert with the weekend indicator (with beginner terrain as the omitted category). These regressions find slightly positive but insignificant coefficients for intermediate and advanced terrain, and a positive and significant coefficient for expert terrain. A regression including an indicator variable for any expert terrain and a continuous variable for the share of expert terrain suggests that the weekend effect is mostly associated with the former. This is consistent with resorts needing expert terrain to be in expert skiers' choice sets, but with the exact amount being less important (or imprecisely measured).

[30] Guarantees, when offered, typically offer the skier a 100 percent refund on that's day lift ticket so long as the skier does not take more than one run.

[31] We define two resorts as competitors if they are within 50 miles (as the crow flies), are not under common ownership, and either both or neither have expert terrain. Permuting the definition does not change the results, nor does changing the mileage filter or the functional form of population.

Table 4—Variation in Weekend Effect by Resort Characteristics

Dependent variable: Resort-reported snow	(1)	(2)	(3)	(4)
Weekend (Sat. and Sun.)	0.175*	0.044	−0.370**	−0.365*
	(0.091)	(0.083)	(0.188)	(0.187)
Interaction effects with weekend				
Expert terrain > 0		0.134	0.142*	0.141*
		(0.084)	(0.078)	(0.076)
No money back guarantee		0.175*	0.144*	0.147*
		(0.097)	(0.077)	(0.076)
Ln(population within 150 miles), normalized			0.206*	0.202*
			(0.106)	(0.108)
1/(1 + number of competitors within 50 miles)			0.677*	0.697*
			(0.395)	(0.405)
West			0.480	0.478
			(0.302)	(0.306)
Publicly traded owner				−0.044
				(0.263)
Government owner				−0.115
				(0.121)
Observations	39,920	39,920	39,920	39,920
Unique days	692	692	692	692
R^2	0.341	0.342	0.342	0.342

Notes: Each specification is estimated using OLS regressions with a control for weighted SNODAS snow and fixed effects for weeks (Wed.–Tues.) and resort. We consider resorts competitors if they are within 50 miles, are not under common ownership, and either both or neither have expert terrain. Standard errors allow for clustering within both day and resort.

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

of the many cases with less than one inch of SNODAS snow. So instead Table 5 reports simple weekend/weekday ratios for resort reports, government reports, and resort/government (i.e., a ratio of weekend/weekday ratios). Consistent with our regression results, Table 5 shows a positive weekend/weekday ratio of ratios for each of the four key groups of resorts—expert, nonexpert, guarantee, nonguarantee—with the largest for expert and nonguarantee. We get similar results using weather station snow instead of SNODAS (online Appendix Table 7).

E. *Postscript: The Rise of Cheap Monitoring Technology and the Fall of the Weekend Effect?*

Finally, we explore the effects of changes to the information environment on exaggeration. On January 8, 2009, *Skireport.com* introduced a feature in its popular iPhone application that makes it easier for users to file "first-hand" reports, in real time (e.g., from the chair lift). These reports are then posted below the resort's official snow report. Online Appendix Table A8 shows that exaggeration disappears following the app launch, with some evidence that effect is driven by resorts with better iPhone reception. We caution that these results are merely suggestive, since there may be other differences between resorts with better and worse iPhone reception, and/or other contemporaneous technology shocks (e.g., slopeside webcams).

Table 5—Average Resort-Reported and Government-Reported Snowfall by Weekday/Weekend and Resort Characteristic

	Weekday snow	Weekend snow	Weekend/Weekday ratio
Expert terrain resorts			
Resort reports	1.46	1.82	1.25
SNODAS	1.21	1.31	1.08
Resort/SNODAS	1.20	1.39	**1.15**
N resort-days	13,139	5,307	
No expert terrain resorts			
Resort reports	1.20	1.46	1.22
SNODAS	1.14	1.27	1.12
Resort/SNODAS	1.06	1.15	**1.09**
N resort-days	15,169	6,305	
No guarantee resorts			
Resort reports	1.41	1.82	1.29
SNODAS	1.21	1.31	1.09
Resort/SNODAS	1.17	1.38	**1.19**
N resort-days	11,262	4,551	
Money-back guarantee resorts			
Resort reports	1.26	1.50	1.19
SNODAS	1.15	1.27	1.11
Resort/SNODAS	1.10	1.18	**1.07**
N resort-days	17,046	7,061	

Note: The figures in bold are the ratio of the ratios (i.e. the ratio of the resort-to-SNODAS ratio on weekends to that on weekdays).

III. Conclusion

We provide some unusually sharp empirical evidence on the extent, mechanics, and dynamics of deceptive advertising.

Ski resorts self-report significantly and substantially more natural snowfall than comparable government reports. Resorts also report more snow on weekends than weekdays, even though there is little evidence from climatology or government measures that snowfall varies across the week. There is some evidence that resorts with greater benefits from exaggerating—those with expert terrain and those without money-back guarantees—do it more. The evidence on whether competition restrains or encourages exaggeration is inconclusive. In all, the results suggest that deceptive advertising about product quality varies sharply with incentives, both within resorts (over time, at high-frequencies), and across resorts.

Although our setting may be unusual with its high-frequency variation in product quality, we speculate that our findings are broadly applicable. They relate to many classes of models on signaling, deception, obfuscation, and search costs. Search and information costs loom large in many other markets where product availability and pricing vary at high frequencies. Some of these markets presumably have conditions that are even more ripe for deceptive advertising than ours, with, for example, purchase decisions that are lower-stakes, have quality realizations with longer lags, or embed fewer opportunities for repeat play and learning.

A particularly important direction for future research is to combine evidence on the nature and dynamics of deceptive advertising with richer evidence on consumer

responses. This is critical for examining whether and how consumers pierce the veil of deception, and for measuring the welfare effects of deceptive advertising and innovations that amplify or discourage it.

REFERENCES

Beales, Howard, Richard Craswell, and Steven C. Salop. 1981. "The Efficient Regulation of Consumer Information." *Journal of Law and Economics* 24 (3): 491–539.

Cabral, Luís, and Ali Hortaçsu. 2010. "The Dynamics of Seller Reputation: Evidence from eBay." *Journal of Industrial Economics* 58 (1): 54–77.

Cerveney, R., and Robert C. Balling, Jr. 1998. "Weekly Cycles of Air Pollution, Precipitation, and Tropical Cyclones in the Coastal NW Atlantic Region." *Nature* 394 (6693): 561–63.

Chen, Yuyu, Ginger Zhe Jin, Naresh Kumar, and Guang Shi. 2012. "Gaming in Air Pollution Data? Lessons from China." *B. E. Journal of Economic Analysis and Policy* 13 (3): 1–43.

Dranove, David, and Ginger Zhe Jin. 2010. "Quality Disclosure and Certification: Theory and Practice." *Journal of Economic Literature* 48 (4): 935–63.

Ellison, Glenn, and Sara Fisher Ellison. 2009. "Search, Obfuscation, and Price Elasticities on the Internet." *Econometrica* 77 (2): 427–52.

Ettinger, David, and Phillipe Jehiel. 2010. "A Theory of Deception." *American Economic Journal: Microeconomics* 2 (1): 1–20.

Forbes, Silke J., Mara Lederman, and Trevor Tombe. 2015. "Quality Disclosure Programs and Internal Organizational Practices: Evidence from Airline Flight Delays." *American Economic Journal: Microeconomics* 7 (2): 1–26.

Forster, P., and S. Solomon. 2003. "Observations of a 'Weekend Effect' in Diurnal Temperature Range." *Proceedings of the National Academy of Science* 100: 11225–30.

Fry, John. 2008. "Lift Ticket Price Sets Record $92." *Skiing Heritage Journal* 20 (1): 1513.

Glaeser, Edward L., and Gergely Ujhelyi. 2010. "Regualting Misinformation." *Journal of Public Economics* 94 (3–4): 247–57.

Holmstrom, Bengt. 1999. "Managerial Incentive Problems: A Dynamic Perspective." *Review of Economic Studies* 66 (1): 169–82.

Hubbard, Thomas N. 2002. "How Do Consumers Motivate Experts? Reputational Incentives in an Auto Repair Market." *Journal of Law and Economics* 45 (2): 437–68.

Jin, Ginger Zhe, and Andrew Kato. 2006. "Price, Quality, and Reputation: Evidence from an Online Field Experiment." *RAND Journal of Economics* 37 (4): 983–1005.

Jolls, Christine, and Cass Sunstein. 2006. "Debiasing through Law." *Journal of Legal Studies* 35 (1): 199–241.

Kartik, Navin, Marco Ottaviani, and Francesco Squintani. 2007. "Credulity, Lies, and Costly Talk." *Journal of Economic Theory* 134 (1): 93–116.

Lazear, Edward P. 1995. "Bait and Switch." *Journal of Political Economy* 103 (4): 813–30.

Luca, Michael. 2011. "Reviews, Reputation, and Revenue: The Case of Yelp.com." Harvard Business School (HBS) Working Paper 12-016.

Luca, Michael, and Jonathan Smith. 2015. "Strategic Disclosure: The Case of Business School Rankings." *Journal of Economic Behavior and Organization* 112 (C): 17–25.

Luca, Michael, and Georgios Zervas. 2013. "Fake It Till You Make It: Reputation, Competition, and Yelp Review Fraud." http://people.hbs.edu/mluca/fakeittillyoumakeit.pdf.

Mayzlin, Dina, Yaniv Dover, and Judith Chevalier. 2014. "Promotional Reviews: An Empirical Investigation of Online Review Manipulation." *American Economic Review* 104 (8): 2421–55.

Mullainathan, Sendhil, Joshua Schwartzstein, and Andre Shleifer. 2008. "Coarse Thinking and Persuasion." *Quarterly Journal of Economics* 123 (2): 577–619.

Petersen, Mitchell A. 2009. "Estimating Standard Errors in Finance Panel Data Sets: Comparing Approaches." *Review of Financial Studies* 22 (1): 435–80.

Rabe-Hesketh, Sophia, Anders Skrondal, and Andrew Pickles. 2004. "Generalized Multilevel Structural Equation Modeling." *Psychometrika* 69 (2): 167–90.

Sauer, Raymond D., and Keith B. Leffler. 1990. "Did the Federal Trade Commission's Advertising Substantiation Program Promote More Credible Advertising?" *American Economic Review* 80 (1): 191–203.

Zinman, Jonathan, and Eric Zitzewitz. 2016. "Wintertime for Deceptive Advertising?: Dataset." *American Economic Journal: Applied Economics.* http://dx.doi.org/10.1257/app20130346.