URI SIMONSOHN*

In 2007, *Consumer Reports* released, and two weeks later retracted, a flawed report on the safety of infant car seats. Analyzing data from 5471 online auctions for car seats ending before, during, and after the information was considered valid, this article shows that (1) consumers responded to the new information and, more surprisingly, (2) they promptly ceased to do so when it was retracted. Because of the random nature of the flawed ratings, this first finding demonstrates that expert advice has a causal effect on consumer demand. The second finding suggests that people's inability to willfully ignore information is not as extreme as the experimental evidence in the psychological literature suggests.

Keywords: expert advice, field data, invalid information, hindsight bias, infant car seats

Lessons from an "Oops" at *Consumer Reports*: Consumers Follow Experts and Ignore Invalid Information

On January 4, 2007, *Consumer Reports* (*CR*) announced its newest safety assessment of infant car seats. The new information, accompanied by the striking image of a car seat flying loose inside a car, differed dramatically from its previous assessment published two years earlier: The correlation of the 2005 and 2007 rankings was r = –.08. Two weeks after the release, *CR* retracted the new safety assessment because of flaws in how the crash simulations were conducted. In short, *CR*, arguably the most trusted source of information for consumer purchases, created new information that was uncorrelated with existing beliefs and then credibly deemed such information invalid shortly afterward.

This article studies the consequences of the release and later retraction of new car seat safety information by analyzing a data set of online auctions for car seats taking place before, during, and after the information was considered valid. The results indicate that consumers promptly responded to both the initial information shock and its retraction. For every position lost/gained in the safety ranking *CR* provided, the average car seat price dropped/increased 3%. (The average car seat moved four positions in the ranking.) More significant, just a few days after the retraction, these price changes were completely eliminated. In other words, the invalidated information no longer affected consumer behavior.

These results directly speak to two important marketing questions. The first is whether expert advice sways consumer demand. This question is typically difficult to answer convincingly because expert recommendations may be correlated with other information consumers hold; therefore, an association between what experts recommend and what consumers do should not be interpreted causally. That *CR* provided wrong information that was uncorrelated with existing beliefs created a unique opportunity to study this question. The large impact of the new safety assessment on auction prices provides direct evidence that experts causally influence consumer demand.

The second question is whether consumers are able to ignore information after it is revealed to be incorrect, something marketing practitioners have often wished their consumers would do. For example, companies are often interested in eliminating negative associations with their brand that have resulted from behaviors in which the company no longer engages, in dismissing rumors about their products, and in eliminating beliefs about their products or services that have been contradicted by scientific evidence.

*Uri Simonsohn is Assistant Professor of Operations and Information Management, the Wharton School, University of Pennsylvania (e-mail: uws@wharton.upenn.edu). The author thanks On Amir, Gerard Cachon, Uri Gneezy, George Loewenstein, Devin Pope, Justin Sydnor, and the anonymous *JMR* reviewers for valuable comments and suggestions. All remaining errors are the author's responsibility. Duncan Simester served as associate editor for this article.

ISSN: 0022-2437 (print), 1547-7193 (electronic)

Both the lay public and marketing researchers commonly believe that people are unable to ignore information after they have received it. A large experimental literature, briefly reviewed in the next section, provides evidence that supports such belief. The findings in this article, the first examination of this question outside the lab, provide a strong counterexample.

EXISTING EVIDENCE ON ADVICE AND CONSUMER DEMAND

In studying the impact of consumer advice on consumer demand, this article is related to recent marketing literature that has studied the impact of word of mouth on demand (e.g., Boatwright, Basuroy, and Kamakura 2007; Chevalier and Mayzlin 2006; Godes and Mayzlin 2004, 2009) and economics literature that has studied the impact of providing expert quality assessment on consumer choice (Cutler, Huckman, and Landrum 2004; Jin and Leslie 2003; Leemore and Dranove 2005; Pope 2009).

In comparison with these literature streams, the current study provides a more fine-grained dependent variable and a simple identification strategy. The dependent variable used here consists of individual auction prices and individual bid amounts rather than aggregate market share or sales. In turn, the causal effect of the advice is identified by virtue of its having been erroneous and uncorrelated with existing information. Previous research attempts have identified advice based on discontinuities of the variable on which the advice is based or the timing of when the information is publicly released.

The current research is also related to Freedman, Schettini, and Lederman's (2009) study, which examines the consequences of product recalls on demand for toys. However, because recalled toys are not available for sale, their data cannot answer the questions of interest herein. Because the car seat data do not include items unaffected by *CR*'s safety assessment, they cannot be used to study the spillover effects in which Freedman, Schettini, and Lederman are interested. Thus, the two efforts are complementary with virtually no overlap.

EXPERIMENTAL EVIDENCE ON IGNORING INFORMATION

An overwhelming body of experimental evidence suggests that people cannot voluntarily ignore information they possess. For example, previous work on the debriefing paradigm has shown that experiment participants receiving false feedback continued to be influenced by it after the experimenters retracted it (Anderson, Lepper, and Ross 1980; Ross, Lepper, and Hubbard 1975). Hundreds of studies have demonstrated that numerical estimates are influenced by starting points (anchors), even when these are transparently irrelevant (and thus should obviously be ignored), such as those obtained through a roulette wheel (Tversky and Kahneman 1974). Hundreds of studies have also documented the hindsight bias, in which people fail to ignore outcomes they have just learned when predicting the beliefs of uninformed people (Fischhoff 1975). A few dozen studies have demonstrated the dilution effect, in which the presence of irrelevant information reduces the weight placed on relevant information even when respondents are asked to judge the relevance of information before making a judgment (Nisbett, Zukier, and Lemley 1981). Finally, a large literature has demonstrated that (mock) juries do not successfully follow instructions to ignore inadmissible evidence (for a meta-analysis, see Steblay et al. 2006).

A noteworthy finding across these and related literature streams is that people fail to ignore information even when forewarned that it is invalid (Gilbert, Tafarodi, and Malone 1993). For example, in the original demonstration of the hindsight bias, Fischhoff (1975) used the same stimuli (though different respondent populations) in experiments in which respondents were and were not forewarned that they would receive to-be-ignored information. The rate of hindsight bias across both experiments was almost identical: The respondents exaggerated the probability of events whose outcome was known by 36% and 32% percentage points, respectively (see Experiments 2 and 3). Randomizing forewarning across conditions of the same experiment, Schul (1993) finds that forewarning leads to a nonsignificant increase in the response to invalid information. Although the experimental evidence that people cannot ignore invalid information is overwhelming, the magnitude of this phenomenon in decisions or purchases outside the laboratory has not been assessed.

THE NEW INFORMATION AND ITS RETRACTION

Consumer Reports publishes a monthly magazine that provides "expert, independent and not-for-profit" reviews aimed at aiding consumers' purchase decisions. Product reviews are typically summarized through rankings based on an overall index of quality and/or safety.

Car seats is one of the product lines *CR* routinely reviews. The February 2007 issue (released to the press on January 4) included a new ranking for 11 infant car seat models. Its previous evaluation, for 8 models, had been published in the May 2005 issue (for both rankings, see Figure 1).

For this new ranking, *CR* modified the testing procedure. In addition to the standard test, in which frontal car collisions are simulated at 30 miles per hour (mph), *CR* conducted a second test that supposedly simulated side-impact collisions at 38 mph. However, because of miscommunications with the company conducting the test, the collisions of this second test were simulated at 70 mph. At such speeds, safety depends almost exclusively on a vehicle's structural integrity; therefore, the results are "nearly meaningless" (*CR* 2007, p. 32). The retraction of the new ranking took place two weeks after it had been released, on January 18, 2007; however, the magazine did not make a new evaluation of infant car seats available until October 2007.

The empirical analyses here focus on the six car seat models covered by both the 2005 and the 2007 (retracted) rankings because, for these, the new information can be classified as good or bad news. Four of these models were reviewed more negatively in 2007 than in 2005, and two were reviewed more positively.

One of these six models, the Evenflo Discovery, failed not only the test erroneously conducted at 70 mph but also the regular 30 mph test. This means that the negative information about it remained valid after January 18, 2007. The importance of this is discussed in greater detail subsequently.

The empirical section analyzes short-term changes in bids and closing prices of car seat auctions as a result of the new information *CR* released. Therefore, it is important to

Figure 1
CONSUMER REPORTS CAR SEAT RANKINGS

A: 2005

Key number	Brand & model	Price	Overall score (0–100; P F G VG E)	Test results: Crash protection – LATCH	Crash protection – Belt	Ease of use – LATCH	Ease of use – Belt	Fit to vehicle – LATCH	Fit to vehicle – Belt	Weight (lb.) Seat/base	Features: Easy harness adjuster	Features: Side-impact material
	INFANT SEATS *Typically fit children up to 20 or 22 lb.*											
1	**Britax** Companion	$200		⊙	◓	◓	◓	⊙	⊙	10/8	•	•
2	**Safety 1st** Designer 22 22-325	80		⊙	⊙	◓	◓	⊙	⊙	9/5		
3	**Peg Perego** Primo Viaggio	170		◓	◓	◓	◓	⊙	⊙	8/7	•	•
4	**Graco** SnugRide 8465	80		◓	⊙	◓	◓	⊙	⊙	8/5	•	
5	**Baby Trend** LATCH-LOC Adjustable Back 6169	90		◓	◓	◓	◓	○	◓	9/12	•	•
6	**Evenflo** Discovery 1	50		○	◓	○	○	◓	◓	6/3		
7	**Combi** Tyro 8000	150		○	⊙	◓	○	○	◒	9/7	•	•
8	**Evenflo** PortAbout 5 Premier Comfort Touch 3751324	85		●	○	◓	◓	⊙	⊙	8/5		

B: 2007 (Retracted)

Key number	Brand & model	Price	Overall score (0–100; P F G VG E)	Test results: Crash protection – Belt	Crash protection – LATCH	Ease of use	Fit to vehicle
	INFANT SEATS *Most accommodate children 5 to 22 lb.*						
1	**Baby Trend** Flex-Loc Adjustable Back	$90	64	◓	○	◓	⊙
2	**Graco** SnugRide with EPS	90	61	○	○	○	⊙
3	**Chicco** KeyFit	150	34	○	◒	⊙	⊙
4	**Peg Perego** Primo Viaggio SIP	230	27	◓	●	◓	⊙
5	**Compass** I410	130	26	○	●	◓	◓
6	**Evenflo** Embrace	70	25	◓	●	○	◓
7	**Britax** Companion	190	20	●	●	◓	⊙
8	**Graco** SafeSeat	150	18	●	●	◓	◓
9	**Safety 1st** Designer 1	120	16	●	●	○	◓
10	**Combi** Centre ST	110	12	●	●	◒	◓
	NOT ACCEPTABLE; DID NOT MEET FEDERAL STANDARD						
11	**Evenflo** Discovery	60					
	NOT ACCEPTABLE; COULD NOT BE INSTALLED SECURELY						
12	**Eddie Bauer** Comfort D 1	200					

consider the likely mechanisms by which auction bidders learned of this information to have a sense of (1) how quickly they could be expected to respond to the information and (2) how much they could be expected to know about the retracted information after it was retracted.

Because of natural delays in the physical distribution of the paper copy of *CR*, it is unlikely that it played an important role.[1] Furthermore, the retraction was not published in print until the April issue, though *CR*'s Web site (www.consumerreports.org) promptly included both the new information and the retraction on January 4 and 18, 2007, respectively. However, data from comScore suggest that it is not common for car seat buyers to visit *CR*'s Web site shortly before making a purchase.[2] In *CR*'s sample of Internet users, only 15% of people buying a car seat online in 2007 visited its Web site during the 30 days preceding the purchase, and just 5% did so on the day of the purchase.

In contrast, data from archives of printed and broadcasted news suggest widespread coverage of the information. For example, Newsbank's database of local newspapers includes 51 newspaper articles about the safety ratings published between January 4 and 6 and 116 retraction stories between January 18 and 20. Similarly, LexisNexis's database of transcripts of broadcasted programs from television and radio, including CNN, Fox News, and NPR, contains 790 and 1277 stories, respectively. Finally, the archived pages of major news Web sites at the Internet Archive (http://www.archive.org) show that the story also received ample coverage on the Web. For example, at cnn.com, the car seat story was one of its ten most popular on January 4.

In light of all this, it seems likely that the majority of car seat buyers learned of the information of interest from the news coverage it received rather than from *CR* directly. This is important for two reasons: First, it validates conducting the empirical analyses on short-term fluctuations in valuations, and second, it reduces the concern that people buying a car seat after the retraction were never aware of the retracted information. On the one hand, the multiple news sources that so heavily covered the *CR* story continued to make that information available after the retraction; on the other hand, parents a few days away from buying a car seat are likely to have paid attention to headlines referring to car seat safety.

THE AUCTIONS DATA

Source of the Data

The data were purchased from Advanced Economic Research Systems, an official data service provider for the largest online auction Web site.[3] The data set includes all auctions listed in the infant car seat category between three months before and three months after the information was released. As mentioned previously, the empirical analyses focus on auctions for the six car seat models covered by both the 2005 and the 2007 *CR* rankings (N = 5471).

Descriptive Statistics

Table 1 presents descriptive statistics for several auction-level variables, tabulated by car seat model. It also reprints the ranking and overall score for the 2005 and 2007 *CR* car seat reviews. Auctions for different car seat models differ substantially in terms of the number of observations in the sample and both their starting and final prices. To include all models into the same regressions meaningfully, dollar variables are deflated by the average selling price of the corresponding car seat model so that the regression estimates correspond to percent changes in prices rather than dollar changes. (As discussed subsequently, the results are robust to various alternative operationalizations of the dependent variable.)

The auction site offers sellers the option to pay additional fees to add special features to their listing (e.g., an additional photo, bold font). The total count of all such features

[1]For example, the library at the university where I work did not receive its copy of *CR* until January 11, 2007, a full week after the information was released.

[2]ComScore has a panel of two million Internet users. Data from a sample of 100,000 are available for download from http://wrds.wharton.upenn.edu.

[3]Because of a nondisclosure agreement, the online auction Web site remains anonymous.

Table 1
DESCRIPTIVE STATISTICS BY CAR SEAT MODEL

	Brand *Model*					
	Britax *Companion*	***Safety 1st*** *Designer*	***Peg Perego*** *Viaggio*	***Graco*** *Snug Ride*	***Baby Trend*** *Adjustable Back*	***Evenflo*** *Discovery*
Rankings (rating) in Consumer Reports						
In 2005	1st (90)	2nd (88)	3rd (70)	4th (69)	5th (55)	6th (45)
In 2007 (retracted)	7th (20)	9th (16)	4th (27)	2nd (61)	1st (64)	11th (0)
Means (Standard Deviation) for Key Variables						
Number of observations	606	243	1327	2682	312	301
Final price (sold items)	\$99.5	\$33.5	\$87.1	\$40.5	\$56.4	\$15.1
	(28.72)	(17.85)	(50.06)	(27.25)	(26.84)	(10.03)
Shipping (sold items)	\$25.3	\$21.3	\$25.1	\$20.3	\$21.1	\$19.8
	(8.41)	(6.06)	(18.11)	(10.18)	(10.81)	(9.77)
Starting price	\$64.22	\$17.49	\$48.14	\$26.61	\$45.48	\$10.59
	(48.43)	(20.76)	(44.63)	(27.62)	(34.22)	(9.77)
Percentage sold	72.1%	62.1%	71.8%	70.9%	62.1%	60.5%
Number of bids (sold items)	13.19	10.48	13.19	9.98	12.11	7.98
	(8.57)	(7.56)	(10.21)	(7.19)	(8.70)	(5.68)
Number of (paid) extra features included with listing	1.17	1.06	1.14	1.28	1.10	1.12
	(.52)	(.23)	(.84)	(.75)	(.52)	(.74)
Percentage of items known to be new	70.1%	75.3%	33.8%	22.9%	51.6%	25.3%

included in a listing is employed as a control in the price regressions; across all models the average number of paid features is similar. Auctions can be set to last 3, 5, 7, or 10 days; the mean in the sample was 5.2 days, and it does not vary much across car seat models.

The new/used status is known for approximately 75% of listings. For the remaining 25%, the items' descriptions were used to infer it. For 7.5% of listings, this was not possible; thus, the price regressions include both new and used dummies, with "unknown" as the omitted category.

The 5471 auctions in the sample were listed by 2825 different sellers. The 4 sellers with the greatest volume listed 18.7% of all auctions, and 62% of sellers listed a single car seat auction. A total of 9248 bidders participated in these auctions; 72% participated in exactly 1, and only 6% participated in 4 or more. Given the small number of repeated observations per bidder and the endogeneity involved in participating in more than one auction, conducting within-bidder comparisons before and after the information shocks is not practical.

Although bidders are allowed to retract their bids, doing so is rare. Furthermore, the retraction rate was similar in the two weeks before the release of the 2007 *CR* ratings (M = .51%), the two weeks they were considered valid (M = .48%), and the two following weeks (M = .44%; $\chi^2(2) = .13$, $p = .93$).

Key Predictors

For the empirical analyses that follow, a key predictor is an auction's ending date. Comparisons are made for auctions ending before the information was released, while it was considered valid, and after it was retracted. These time periods are referred to as Before, During, and After. (Before was typically the omitted category in the regressions.)

Some of the analyses use the change in a car seat's rankings as the predictor of interest (interacted with the time dummies). The ranking change variable, ΔRanking, is defined as a model's ranking in 2007 less its ranking in 2005. For example, the car seat Adjustable Back went from being ranked fifth in 2005 to being ranked first in 2007, so ΔRanking is –4 for that model.

EMPIRICAL ANALYSES

This section is divided into four subsections. The first subsection focuses on the impact of the information shocks on final prices. The second subsection discusses and addresses the issue of statistical independence across observations in those analyses. The third subsection expands the analyses to all bids and employs quantile regressions to assess how widespread the findings are among nonwinning bids. The fourth subsection examines the potential role of sellers. Given that a natural experiment is behind the data-generating process, all analyses are reduced form; structural estimations are not needed to identify the effects of interest.

Impact of Information Shocks on Final Prices

The impact on final prices is assessed with two complementary approaches. The first employs a car seat's change in the ranking between 2005 and 2007 (ΔRanking) to predict changes in selling prices. The second estimates price changes for each car seat model separately, relaxing the functional form assumption of a linear impact of ranking change on percent price change.

The first set of results are from regressions that have auctions as the unit of analysis and selling price as the dependent variable (deflated by average prices), and the predictors of interest are the interactions of the time dummies (Before, During, and After) with ΔRanking. More specifically, the following equation for auction i of car seat model k was estimated:

$$(1)\ \mathrm{Price}_{i,k}/\mathrm{AvgPrice}_k = \beta_0 + \beta_1 \times \mathrm{During}_i + \beta_2 \times \mathrm{After}_i + \beta_3 \times \Delta\mathrm{Ranking}_k + \beta_4 \times \mathrm{During}_i \times \Delta\mathrm{Ranking}_k + \beta_5 \times \mathrm{After}_i \times \Delta\mathrm{Ranking}_k + \gamma \times \mathrm{controls}_{i,k},$$

where AvgPrice_k is the average selling price of car seat model k.

If bidders responded to information while it was valid, $\beta_4 < 0$ (indicating that prices are lower when a car seat drops in the ranking). If people stop responding to information after it is retracted, $\beta_5 = 0$.

Because ΔRanking varies at the car seat model level rather than the auction level (note the k subscripts), standard errors are clustered by time × car seat model, leading to 18 clusters. The regressions were estimated including observations from auctions ending between three weeks before the new ranking was released and three weeks after it was retracted to focus on variation from the relevant period. For robustness, these regressions were also estimated for larger time windows (five weeks and three months for both Before and After periods).

Table 2 presents the results. Columns 1–3 show ordinary least squares (OLS) results for auctions receiving at least one bid, and Columns 4–8 present results from censored normal regressions, which include all auctions, treating those with 0 bids as censored at their starting price.[4] The point estimates of interest are displayed in bold.

Column 1 presents the base specification, controlling only for item attributes. β_4 is –.028, indicating that for every position lost in the ranking, the estimated price change during the two weeks in which such ranking was considered valid was a drop of 2.8%. (On average, car seats moved four positions, so the average effect is approximately an 11% price change.) In other words, consumers responded to the new information *CR* provided.

In contrast, β_5 is close to and not statistically different from 0 (β_5 = –.001, p = .93), indicating that after the new ranking was retracted, it no longer influenced auction prices. Columns 2 and 3 add controls for auction-level attributes and competition. The resulting point estimates of interest are similar to those of Column 1.

Column 4 reports the results from a censored normal regression, in which auctions with no bids are treated as left censored (at the starting price) and those sold with the buy-it-now option are treated as right censored (at the buy-it-now price). Importantly, shipping charges are unknown for unsold auctions, and thus neither the dependent variable nor the censoring point includes them. Therefore, the censored normal regression improves on one shortcoming from the OLS regression (selection into sale) but introduces a new one (measurement error in the dependent variable).

Qualitatively, the censored normal results are aligned with the OLS ones: $\beta_4 < 0$ and $\beta_5 \sim 0$. The results from Columns 5 and 6 are similar to those from Column 4, indicating that the point estimates are robust with respect to the time windows employed.

Columns 7 and 8 estimate the regressions separately for new and used car seats. The effect of the information is expected to be larger for new car seats because people are often advised not to buy used ones; therefore, consumers buying used car seats are expected to be less prone to following advice in general and *CR*'s in particular.[5] The effect of the new information is estimated to be much larger and is only statistically significant for new car seats.

A variety of robustness tests were performed on the previous results:

1. *Clustering*. Rather than the errors being clustered at the model-time level (18 clusters), they were clustered at the model level only (6 clusters) and not at all (standard OLS errors). For every column in Table 2, the standard errors for β_4 were largest when clustering by model-time.
2. *Logs*. All columns from Table 2 were reestimated using the log of price as the dependent variable. The point estimates for β_4 were similar: approximately –3% in Columns 1–3 and approximately –7% in Columns 4–6; for new car seats, it was –8.4%, and for used ones, it was –2.1%. All the estimates were significant at the 1% level except the used car seats estimate, which was not significant at the 10% level. None of the point estimates for β_5 were significant at the 10% level.
3. *Out-of-sample average price.* Rather than the dependent variable being computed as the ratio of the selling price to the average selling price of the car seat model during the whole sample, it was computed with the average selling price up to four weeks before the information was released. For Columns 1–4, this means that the dependent variable was computed with a holdout sample. The correlation between the original dependent variable and the one based on previous prices is extremely high (r = .99), so it is not surprising that the resulting point estimates barely changed: β_4 is –.029, –.024, –.026, and –.034 in Columns 1–4, respectively, all significant at the 5% level, and β_5 is .000, .004, .002, and –.013, respectively, none significant at the 10% level.

How Fast Did Consumers Respond?

The regressions from Table 2 make discrete comparisons of Before with During and After. To assess how quickly consumers responded to the release and retraction of information, regressions with a more fine-grained measure of time are necessary. To this end, time was split into three-day intervals, and the interaction of the resulting three-day dummies with ΔRanking was used to study price dynamics more closely. (The sample is not dense enough to allow an even finer measure of time.)

Figure 2 plots the point estimates of these regressions. Auctions ending more than six days before the release of the new ranking are the omitted category. The figure has three key features. First, it shows no apparent trend before the new information is released (dummies for days –1 to –3 and –4 to –6 are near 0). Second, it shows a drop of nearly 4%, per position lost in the ranking, immediately after the release of the information stabilizing around 3% for the remaining During days. Third, it shows that by day 4 or 5 after the retraction, the effect of the invalid information was fully eliminated.

Prices at the Individual Car Seat Level

The previous analyses impose an arbitrary functional form on the relationship between information and prices—

[4]Censored normal regressions have the same basic structure as a Tobit, but they allow each observation to be censored at a different point. This specification is needed here because the starting price differs across observations (e.g., Wooldridge 2001).

[5]For example, CarSeatReviews.org writes, "There are several things you should consider when choosing your toddler car seat. The first is that it should be new. Buying a used car seat is never the best option" (http://www.carseatreviews.org/how-to-shop-for-toddler-car-seats.html, last retrieved March 13, 2010). See also http://blogs.consumerreports.org/cars/2009/06/what-to-do-with-a-used-car-seat.html.

Table 2

REGRESSIONS PREDICTING FINAL PRICE OF CAR SEAT AUCTIONS

	OLS			*Censored Regression Models*				
	Base Specification	*Adds Auction-Design Controls*	*Adds Controls for Competition*	*Includes Unsold Items*	*Adds More Weeks*	*Full Sample*	*Full Sample New Items*	*Full Sample Used Items*
Sample period (in weeks)	3 Before 2 During 3 After	3 Before 2 During 3 After	3 Before 2 During 3 After	3 Before 2 During 3 After	5 Before 2 During 5 After	14 Before 2 During 13 After	14 Before 2 During 13 After	14 Before 2 During 13 After
During (1 = yes, 0 = no): Does auction end during the two weeks in which new ranking was valid?	.159*** (.034)	.145*** (.030)	.150*** (.032)	.180*** (.033)	.159*** (.033)	.169*** (.040)	.238*** (.083)	.128*** (.034)
After (1 = yes, 0 = no): Does auction end after the retraction?	.063** (.029)	.081*** (.027)	.092*** (.031)	.096** (.043)	.081* (.044)	.042 (.048)	.056 (.077)	.042 (.036)
ΔRanking (2007 rating – 2005 rating): Difference in *CR*'s safety of auctioned car seat model	–.033***	–.022*	–.020	–.016*	–.009	–.001	.005	–.003
During × ΔRanking	**–.028** (.011)**	**–.024** (.011)**	**–.026** (.011)**	**–.031*** (.010)**	**–.035*** (.010)**	**–.034*** (.012)**	**–.047*** (.017)**	**–.011 (.014)**
After × ΔRanking	**–.001 (.011)**	**.003 (.011)**	**.001 (.011)**	**–.007 (.010)**	**–.008 (.010)**	**–.006 (.011)**	**–.010 (.014)**	**–.009 (.012)**
Is item known to be new? (1 = yes, 0 = no)	.515*** (.083)	.373*** (.062)	.390*** (.063)	.376*** (.086)	.325*** (.069)	.320*** (.045)		
Is item known to be used? (1 = yes, 0 = no) Excluded category: unknown status	–.048 (.048)	–.033 (.027)	–.028 (.025)	–.067 (.054)	–.098** (.041)	–.042 (.038)		
Includes additional base? (1 = yes, 0 = no)	.131** (.056)	.142*** (.045)	.139*** (.043)	.061 (.047)	.028 (.035)	.022 (.026)	–.243** (.111)	.031 (.030)
Auction's starting price (as percentage of model's average price)		.491*** (.076)	.491*** (.076)	.366*** (.076)	.388*** (.075)	.391*** (.073)	.450*** (.069)	.572*** (.031)
Number of (paid) extra features included with listing		.036*** (.010)	.036*** (.009)	.033*** (.012)	.032** (.013)	.024* (.015)	.005 (.035)	.023 (.015)
Reserve price present (1 = yes, 0 = no)		.243*** (.025)	.243*** (.026)	.104*** (.037)	.147*** (.031)	.154*** (.020)	.282*** (.075)	.122*** (.016)
Offered buy-it-now (1 = yes, 0 = no)		.007 (.032)	.006 (.032)	.126*** (.031)	.101*** (.029)	.099*** (.021)	–.010 (.032)	.106*** (.024)
Number of same-model auctions that day			–.002 (.002)	–.003 (.003)	–.001 (.004)	.003 (.003)	.007 (.006)	–.000 (.002)
Percentage of same-model auctions that day that are new			–.058 (.080)	.023 (.065)	.020 (.069)	–.033 (.079)	–.163* (.097)	.130** (.057)
Number of observations	1052	1052	1052	1438	2227	5471	2480	3346
R^2/pseudo-R^2	.32	.47	.47	.30	.29	.26	.16	.16

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

Notes: The key predictors, indicated in boldface, are the interaction of the change in the *Consumer Reports* (*CR*) ranking, for the corresponding car seat model between the 2005 and 2007 ranking, and the time dummies indicating whether an auction ended during the period when ranking was considered valid or after it. The interpretation of the point estimate is the percentage drop in final prices for auctions of a car seat model that dropped one position in the ranking, during the time period indicated by the dummy variable. Standard errors are clustered by car seat model × time dummy (18 clusters).

Figure 2

DYNAMICS OF AUCTION PRICES FOLLOWING NEW INFORMATION AND ITS RETRACTION

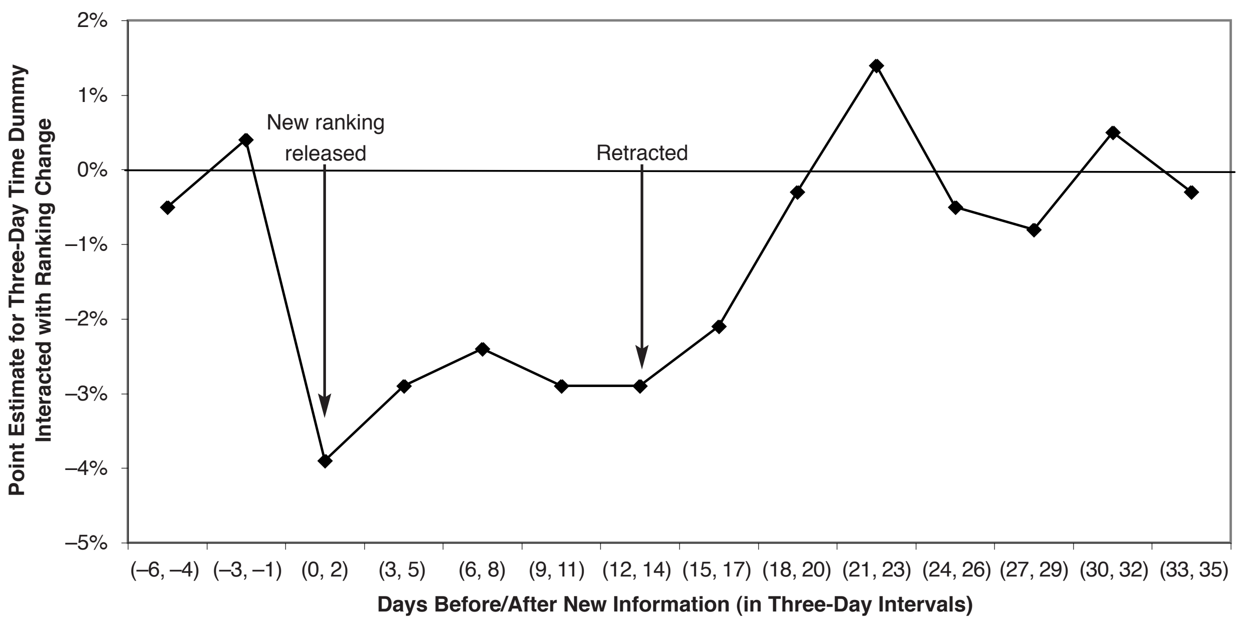

Notes: This figure plots regression point estimates for the interactions between time dummies covering three-day periods and the change in ranking for the corresponding car seat model between 2005 and 2007 (ΔRanking). The point estimate for the first two dummies is almost exactly 0, indicating that there was no price change on those days, predicted with the upcoming change in *Consumer Report* ranking. The information was released on January 4, day 0 in the figure, and retracted on January 18, day 15 in the figure. The regression employed to obtain the plotted point estimates controls for all covariates listed in Column 3 of Table 2.

namely, that prices change linearly in percentage terms in response to changes in the safety ranking. A way around such assumptions is to estimate the effects on final prices separately for each model.

To this end, a set of regressions that interacted the time dummies with car seat model dummies rather than ΔRanking was estimated. The interactions' point estimates are the change in prices for a given model, net of the average change of the other models. A total of 12 regressions were estimated: a censored normal and an OLS for each car seat model in the sample. Figure 3 plots the estimates for these regressions, averaging the point estimate between the censored normal and OLS for each model. The results are consistent with those obtained with ΔRanking.

In particular, the four car seat models that dropped in the rankings of the flawed report experienced a price drop in the During period, while the two car seats that gained in the ranking saw a price increase. The Baby Trend car seat experienced the largest price change, going from the fifth to the first position in the ranking.[6] This provides further evidence that the new information *CR* released causally influenced consumer behavior.

Figure 3 also shows that most price changes vanish in the After period. The notable exception is Evenflo's car seat. This was the car seat that failed not only the 70 mph test but also the standard 30 mph one, and thus its negative information was not invalidated. Again, this suggests that after the retraction, people still had access to the original information (otherwise, they would not know that the Evenflo was unsafe).

Statistical Independence

A limitation of the current natural experiment is that though it contains several thousand observations, there are only two treatments (or four if viewing receiving positive versus negative information as separate treatments), and these are fixed among all observations within a time period. This general concern with natural experiments is amplified in the current context because the multiple products studied are substitutes and because buyers and sellers can observe the actions of others when making decisions. Although standard errors were clustered in the price regressions to account for this, the number of clusters (18) may have been insufficient to address the problem properly (Bertrand, Duflo, and Mullainathan 2004).

This subsection further addresses the issue of independence by treating two weeks' worth of data as an observation, leading to just 14 observations, and assessing how unusual the two weeks of interest (when the information was considered valid) are in relation to the other two-week periods (biweeks) in the data. Note that by using such a long period as an observation, any remaining lack of independence would reduce rather than enhance the odds of documenting marked changes in prices following the release and subsequent retraction of information and thus lead to a conservative test of the hypothesis of interest.

[6]Given the large price change for this car seat model, the regressions from Table 2 excluding auctions for the Adjustable Back model were reestimated. The point estimate for During × ΔRanking remained negative and highly significant ($\beta_4 = -.023$, $p < .01$).

Figure 3
PRICE CHANGES DURING AND AFTER NEW INFORMATION WAS CONSIDERED VALID (BY CAR SEAT MODEL)

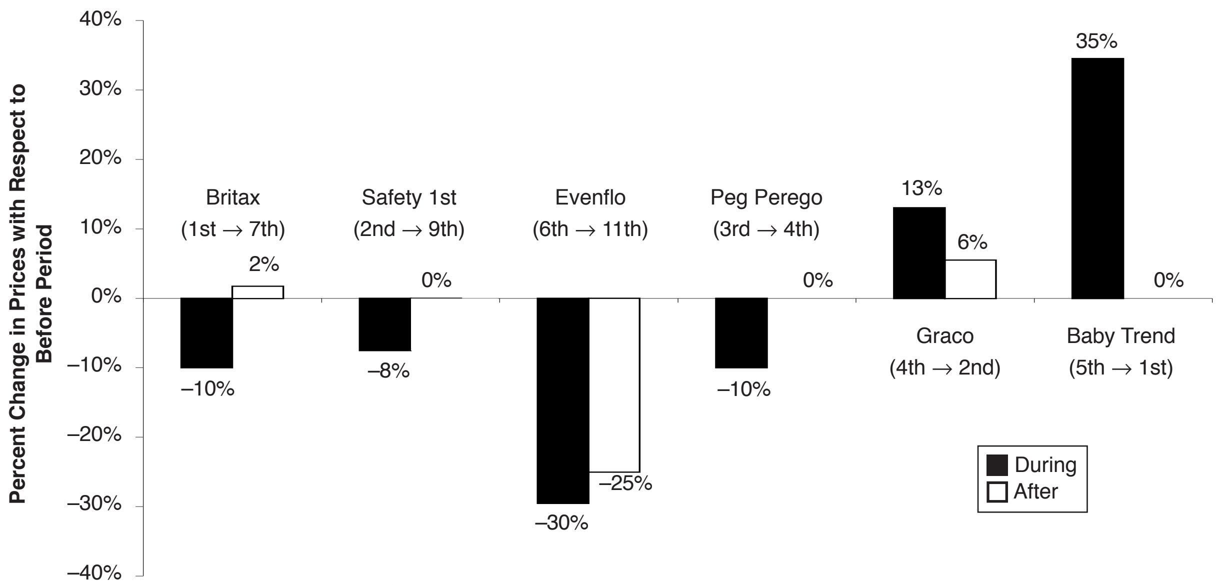

Notes: Labels indicate the brand of the car seat on top and the change in rankings from 2005 to 2007 (in parentheses). During consists of auctions ending while the new information was considered valid; After refers to those ending within three weeks of the retraction. Price changes are with respect to the three weeks before the release of the information. Price changes are average point estimates from censored normal and OLS regressions that control for covariates. See text for description.

Let us begin with the price regressions that employ the interaction of the time dummies with ΔRanking as the key predictors (Table 2). The specification from Column 3 was reestimated, but this time it interacted ΔRanking with biweekly dummies rather than During and After.

Figure 4 plots the results for the roughly six months of data. The figure shows that most biweekly price shocks are noticeably smaller than ±1%. That is, on any given biweek, the average price change in car seat models, as predicted by their future change in the rankings, is less than 1% per position changed. This contrasts with the –2.9% for the biweek of interest.

Considering the average magnitude of price changes across biweeks (M = .67%, SD = .44%), the 2.9% drop is

Figure 4
BIWEEKLY PRICE CHANGES PER POSITION LOST IN NEW RANKING

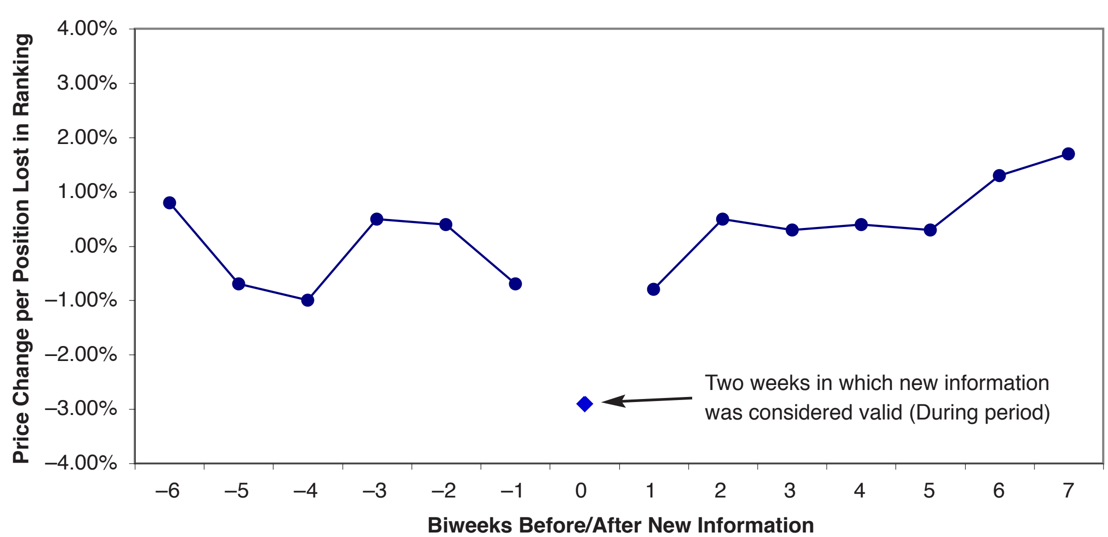

Notes: Each dot in the figure is a point estimate of ΔRanking × biweek dummies from regressions with price as the dependent variable (similar specification to Column 3 in Table 2).

5.3 standard deviations larger than the typical biweekly price change. In other words, after accounting for possible failures of statistical independence, it is still the case that sampling error cannot plausibly account for the response to the information shocks.

Analogous calculations were performed for the regressions that estimate price changes on a car seat model basis (Figure 3). The question here is slightly different. How often is a pattern that is as consistent with the hypothesis that *CR* influenced market prices as that depicted in Figure 3 observed in other biweeks? That is, how often do the four car seat models that dropped in the ranking experience a negative price shock and the other two a positive price shock?

To answer this question, a 1/0 indicator was created for each biweek × car seat estimate. The indicator equals 1 if the car seat experienced a price change in the same direction as that predicted for the During period and 0 if otherwise. Adding this indicator variable across the six car seats for each biweek provides a count of how many models moved in the predicted direction (a sum of 6 indicates that all models had price shocks in the direction predicted for the During period).

Across other biweeks, this variable ranged between 1 and 5 (M = 2.5, SD = 1.24). That is, on an average biweek, 2.5 of the 6 car seats experienced a price shock that matched the direction predicted for the During biweek. Thus, the (perfect) prediction of 6/6 price changes obtained in the During biweek (and in no other biweek) is 2.8 standard deviations more accurate than the typical biweek. This further suggests that the patterns documented previously are unlikely to have been caused by sampling error.

Does Everyone Succeed at Ignoring the Retracted Information?

Although the evidence presented herein suggests that consumers were able to ignore invalidated information, that they did so for every car seat model, and that they did so promptly, the results rely on closing prices and thus depend primarily on auctions' marginal bidders. It is possible that a large portion of consumer fail to ignore the invalidated information, but just not enough of them to influence final prices. This is particularly relevant given that in a lab experiment, Camerer, Loewenstein, and Weber (1989) find that the hindsight bias is attenuated in a market context because those least affected by it trade more.

To address this possibility, regressions similar to those presented previously but with bids rather than auctions as the unit of analysis were estimated. For the purpose of learning about bidders across the full range of valuations, quantile regressions for bids in the 20th, 40th, 60th, and 80th percentiles were estimated. A quantile regression is analogous to a standard linear regression, but instead of maximizing fit for the average value, it does so for a given quantile (see Koenker and Bassett 1978; Koenker and Hallock 2001). Intuitively, rather than assessing how the new information influenced the average bid, it estimates how much it influenced bids around, for example, the 20th percentile of amount bid. Although it is probably incorrect to treat bid amounts as maximum willingness to pay (Zeithammer 2010), it is more reasonable to presume that changes in willingness to pay lead to changes in bid amounts.

Table 3 reports the point estimates and bootstrapped standard errors for these quantile regressions. For all four quantiles examined, During × ΔRanking is estimated as negative and significant, while After × ΔRanking is not statistically different from 0. This is consistent with bidders across the

Table 3

QUANTILE REGRESSIONS WITH BID AMOUNT AS DEPENDENT VARIABLE

	Quantile			
	20%	*40%*	*60%*	*80%*
During (1 = yes, 0 = no): Does auction end during the two weeks in which new ranking was valid?	.049*** (.008)	.060*** (.016)	.039*** (.012)	.059** (.026)
After (1 = yes, 0 = no): Does auction end after the retraction?	.001 (.009)	.013 (.013)	.009 (.013)	.008 (.025)
ΔRanking (2007–2005) Difference in *CR*'s safety of auctioned car seat model	–.001 (.003)	.002 (.003)	.004 (.003)	–.002 (.004)
During × ΔRanking	**–.008** (.003)**	**–.017*** (.003)**	**–.021*** (.003)**	**–.024*** (.006)**
After × ΔRanking	**.001 (.003)**	**.001 (.003)**	**.002 (.003)**	**.007 (.005)**
Is item known to be new? (1 = yes, 0 = no)	.126*** (.020)	.225*** (.021)	.285*** (.030)	.247*** (.049)
Is item known to be used? (1 = yes, 0 = no) Excluded category: unknown status	.024 (.015)	.022 (.017)	.023 (.029)	–.048 (.048)
Includes additional base? (1 = yes, 0 = no)	.017 (.020)	.026 (.019)	.044 (.032)	.096* (.050)
Auction's starting price (as percentage of model's average price)	.897*** (.016)	.817*** (.019)	.718*** (.015)	.586*** (.015)
Number of (paid) extra features included with listing	.014*** (.005)	.017** (.007)	.021*** (.008)	.030*** (.011)
Reserve price present (1 = yes, 0 = no)	.122*** (.022)	.153*** (.022)	.195*** (.029)	.232*** (.054)
Offered buy-it-now (1 = yes, 0 = no)	–.032*** (.011)	–.022 (.015)	–.011 (.015)	.008 (.024)
Number of same-model auctions that day	.001* (.000)	.001*** (.000)	.002*** (.000)	.002*** (.000)
Percentage of same-model auctions that day that are new	.004 (.012)	–.048 (.032)	–.079*** (.029)	–.035 (.038)
Number of observations	4974	4974	4974	4974
Pseudo-R^2	.216	.216	.218	.202

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

Notes: This table shows point estimates from quantile regressions employing bids as the unit of observation and dollar amount of bid as a percentage of average final price for car seat model as the dependent variable. The sample includes auctions between three weeks before January 4 and three weeks after January 18. The key predictors, indicated in boldface, are the interaction between the change in the *Consumer Reports* (*CR*) ranking, for the corresponding car seat model, between the 2005 and 2007 rankings, and the time dummies. Their interpretation is the percentage drop in a bidder's bid amount for a car seat model that dropped one position in the ranking during the period indicated by the dummy variable. Bootstrapped standard errors are in parentheses below point estimates.

range of valuations having responded to the new ranking only while it was valid.

Comparing the point estimates across columns, the effect size of the new information appears larger for the higher quantiles, but this is mechanically true because lower amount bids necessarily have smaller effects (expressed as percentage of final prices). Deflating the point estimates by the bid amount of the corresponding quantile makes the effect sizes much more similar: –1.7%, –2.5%, –2.4%, and –2.0% for the 20th, 40th, 60th, and 80th quantiles, respectively. In short, the evidence is consistent with bidders across the full range of valuations having first responded to the new information and then succeeded at ignoring it after its retraction.

Seller Behavior

Considering that many sellers are individual people in this market (recall that 62% of listings were posted by sellers offering a single auction in the sample), it is possible that at least part of the effects documented for final prices were driven by the actions of sellers (e.g., parents selling their now thought to be unsafe car seats). To assess the potential role of sellers, Figure 5 reports results analogous to those of Figure 4 but for variables under the control of sellers. Figure 5, Panel A, plots the biweekly dynamics of starting prices and number of paid features as a function of the change in the rankings for the car seat model being auctioned.

The figure shows a fairly flat line for starting price, indicating that the magnitude of the biweekly changes in starting prices (M = .69%, SD = .67%) and number of paid features (M = .01, SD = .009) are small. Most important, there is no pattern on the seller side that could account for market price dynamics (e.g., a markedly negative point estimate for the ΔRanking interaction on biweek = 0).

Figure 5, Panel B, shows similar calculations for the number of listings of each car seat model and the percentage of them that are new. There is a minor (not statistically significant) dip in the percentage of car seats that are new offered in biweek 0. This is unlikely to account for the findings for various reasons, including that the price effect is

Figure 5

BIWEEKLY CHANGES

A: Changes in Seller-Controlled Variables per Position Lost in New Ranking

B: Changes in Number of Listings and Percentage of New Listings per Position Lost in New Ranking

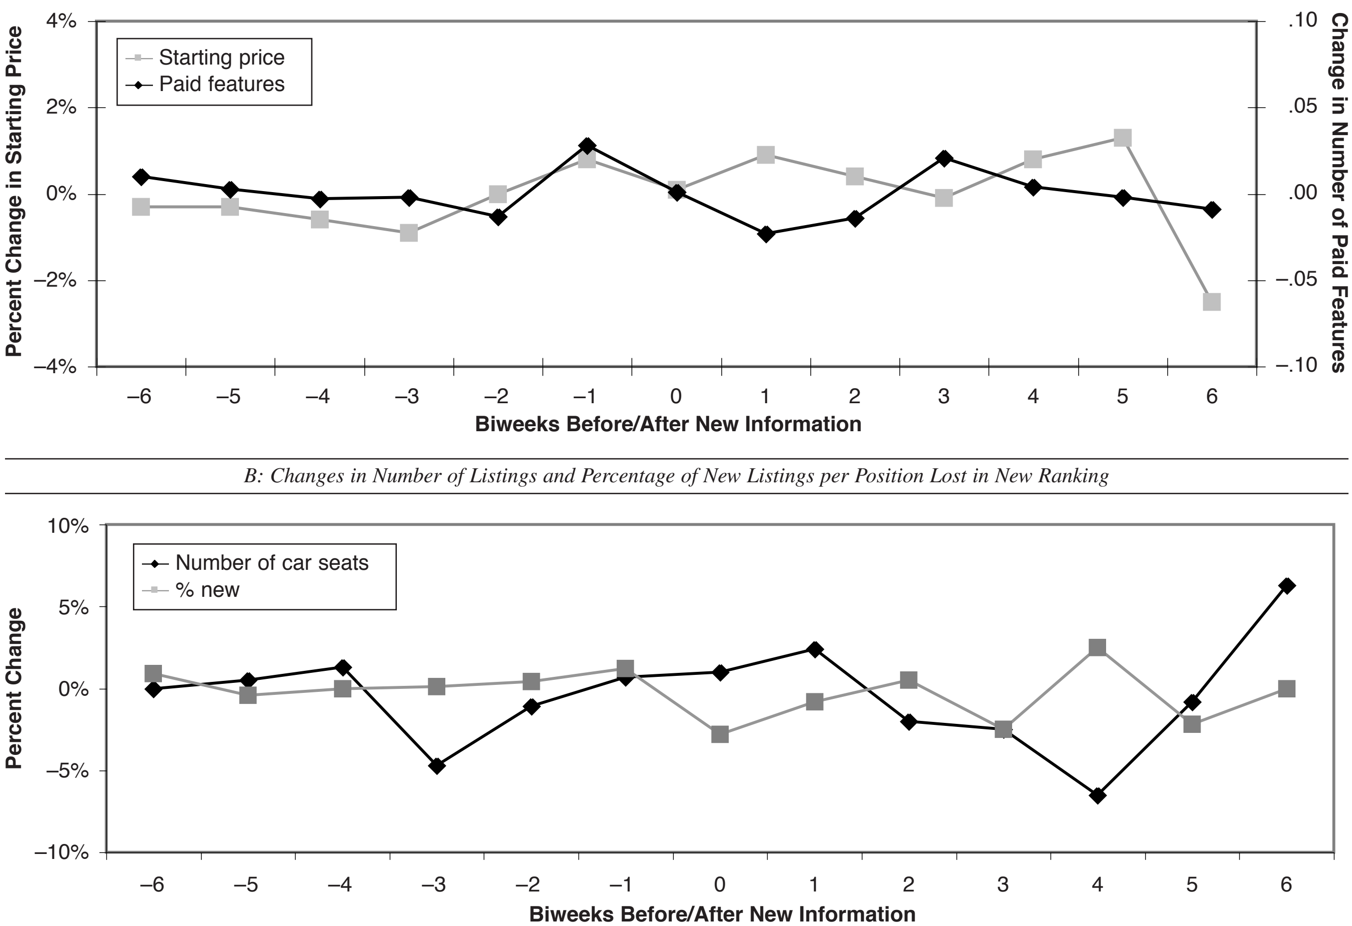

Notes: In Panel A, each dot is a point estimate of ΔRanking × biweek dummies from regressions with starting price and number of paid features as dependent variables and auctions as the unit of observation (N = 5471). In Panel B, each dot is a point estimate of ΔRanking × biweek dummies from regressions, with numbers of biweekly auctions for a given model, and the percentage of these that are new, as dependent variables, with biweek × model as the unit of observation (N = 14 × 6 = 84).

only documented for new car seats and that in biweek = 1 the percentage of new car seats remains lower than normal while prices go back to baseline in biweek = 1. In summary, the evidence suggests that variation in market prices following the new information was primarily, if not exclusively, driven by variation in demand.

This is not too surprising. Individual sellers would probably wait to acquire a new car seat before putting their old one on the market, and they might even have reservations about listing an item for sale they believe to be unsafe. Institutional sellers probably do not have the flexibility to respond to such short-term shifts; had the information remained considered valid for a few months, sellers might have shifted their product lines or revised their prices.

DISCUSSION

Main Findings of the Current Research

This article studies the consequences of an involuntary natural experiment conducted by *CR*, which retracted a new ranking of car seat safety just two weeks after releasing it. The article has two main findings: (1) Consumers promptly responded to the new information, and (2) they successfully ignored the retracted information. Both findings contribute to the understanding of consumer behavior and have clear marketing implications.

The first finding answers perennial questions that marketing researchers and practitioners have great interest in, such as the following: Do third-party reviews causally influence demand? and Do consumers respond to safety information regarding products? These questions are often difficult to answer convincingly because recommendations and safety ratings are not exogenous; rather, they are correlated with other information held by consumers, and thus correlation between the reviews by third parties and demand cannot be interpreted causally. The natural experiment studied herein, in which new information was all but random, provides a rather conclusive answer to these questions: yes.

The second finding, that people were able to discount information fully when it was retracted, also has relevance for both marketing researchers and practitioners. Marketers frequently encounter situations in which they wished consumers would ignore invalid information. For example, companies are often interested in eliminating negative associations with their brand (e.g., "American cars are unreliable"), dismissing rumors about their products (e.g., "The food sold at that restaurant is made from mutant chickens"[7]), and eliminating wrong beliefs that have been debunked by scientific studies (e.g., "Drinking coffee while pregnant leads to fetal development problems"; see Linn et al. 1982).

As mentioned previously, both the lay public and decision-making researchers commonly believe that people are unable to ignore information after they have received it. This study's findings provide a strong counterexample. The next subsection discusses possible explanations for the inconsistent set of results put forward here compared with existing laboratory evidence. Although further research should examine the relative importance of these explanations, the current findings invite researchers to reconsider what the experimental literature predicts regarding the ability of people outside the lab to ignore information.

Why Was Information Ignored Here but Not in the Lab?

A careful read of previous experimental findings sheds some light on potential explanations for why people typically fail to ignore information in the lab but succeeded here. In particular, (1) *CR* providing a substantive and compelling instruction to ignore the retracted information and (2) consumers having access to additional information regarding car seat safety are instantiations of two factors that existing research suggests facilitates successfully ignoring information. In terms of the former, research studying (mock) jury decision making has shown that though juries fail to follow instructions to ignore information when the reason to ignore it is procedural (e.g., an incriminating audiotape obtained without a warrant), they do successfully ignore information when the reason to ignore it is substantive (e.g., an incriminating transcript obtained from a barely audible audiotape) (e.g., Kassin and Sommers 1997). Similarly, recent research on the debriefing paradigm reveals that if respondents are told the feedback they received in the experiment should be ignored because the test used to generate it is fake, they succeed at ignoring such feedback (McFarland, Cheam, and Buehler 2007).

These findings suggest that part of the reason behind respondents' inability to ignore invalid information in experiments is that they do not view the to-be-ignored information as irrelevant as the experimenters believe it should be. That *CR* gave a substantial and compelling reason to ignore the safety information is likely to have contributed to the success consumers had in ignoring it.

Another factor that may have contributed to consumers ignoring the retracted information successfully is that they were able to (and probably did) pursue additional information after learning that the information they had was no longer valid. Consumers desiring to obtain additional information could consult consumer reviews on various Web sites (e.g., Amazon.com), earlier *CR* assessments, safety ratings conducted by other agencies, and so on. Existing research has shown that considering alternative outcomes reduces both the hindsight bias (Arkes et al. 1988; Slovic and Fischhoff 1977) and anchoring (Chapman and Johnson 1999). It has also shown that being able to physically eliminate irrelevant information attenuates the dilution effect (Kemmelmeier 2004). This suggests that in settings in which consumers are free to pursue additional information and dispose of unwanted information more easily, which they typically cannot do in a lab setting but can easily do outside it, they will be more likely to ignore information successfully.

More broadly, the findings emphasize the importance for lab research to set out explicitly to explore the impact of moderators likely to be present outside the lab on findings obtained in the lab. For example, in the context of studying people's ability to ignore information, lab studies have not allowed participants to obtain additional information, postpone the decision, or physically eliminate the irrelevant information. They also have not systematically manipulated (with the exception of the mock jury evidence reviewed previously) the nature of the credibility of the source falsifying the information. Because all these variables differ in the

[7]See http://www.snopes.com/horrors/food/kfc.asp.

field compared with the (existing) lab evidence, carefully manipulating them in further research would be useful to deepen understanding of the external validity of existing lab research on this topic.

REFERENCES

Anderson, Craig A., Mark R. Lepper, and Lee Ross (1980), "Perseverance of Social Theories: The Role of Explanation in the Persistence of Discredited Information," *Journal of Personality and Social Psychology*, 39 (6), 1037–1049.

Arkes, H., D. Faust, T.J. Guilmette, and K. Hart (1988), "Eliminating the Hindsight Bias," *Journal of Applied Psychology*, 73 (2), 305–307.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan (2004), "How Much Should We Trust Differences-in-Differences Estimates?" *Quarterly Journal of Economics*, 119 (1), 249–75.

Boatwright, Peter, Suman Basuroy, and Wagner Kamakura (2007), "Reviewing the Reviewers: The Impact of Individual Film Critics on Box Office Performance," *Quantitative Marketing and Economics*, 5 (4), 401–425.

Camerer, Colin F., George Loewenstein, and Martin Weber (1989), "The Curse of Knowledge in Economic Settings: An Experimental Analysis," *Journal of Political Economy*, 97 (5), 1232–54.

Chapman, Gretchen B. and Eric J. Johnson (1999), "The Limits of Anchoring," *Journal of Behavioral Decision Making*, 7 (4), 223–42.

Chevalier, Judith A. and Dina Mayzlin (2006), "The Effect of Word of Mouth on Sales: Online Book Reviews," *Journal of Marketing Research*, 43 (August), 345–54.

CR (2007), "How Our Car Seat Tests Went Wrong—The Missteps We Made and Some Lessons We Learned," (May), 30–33.

Cutler, David M., Robert S. Huckman, and Mary Beth Landrum (2004), "The Role of Information in Medical Markets: An Analysis of Publicly Reported Outcomes in Cardiac Surgery," *American Economic Review*, 94 (2), 342–46.

Fischhoff, Baruch (1975), "Hindsight Not Equal to Foresight: Effect of Outcome Knowledge on Judgment Under Uncertainty," *Journal of Experimental Psychology: Human Perception and Performance*, 1 (3), 288–99.

Freedman, Seth M., Melissa Kearney Schettini, and Mara Lederman (2009), "Product Recalls, Imperfect Information, and Spillover Effects: Lessons from the Consumer Response to the 2007 Toy Recalls," NBER Working Paper No. 15183, National Bureau of Economic Research (July).

Gilbert, Daniel T., Romin W. Tafarodi, and Patrick S. Malone (1993), "You Can't Not Believe Everything You Read," *Journal of Personality and Social Psychology*, 65 (2), 221–33.

Godes, David and Dina Mayzlin (2004), "Using Online Conversations to Study Word-of-Mouth Communication," *Marketing Science*, 23 (4), 545–60.

——— and ——— (2009), "Firm-Created Word-of-Mouth Communication: Evidence from a Field Test," *Marketing Science*, 28 (4), 721–39.

Jin, Ginger Zhe and Phillip Leslie (2003), "The Effect of Information on Product Quality: Evidence from Restaurant Hygiene Grade Cards," *Quarterly Journal of Economics*, 118 (2), 409–451.

Kassin, Saul M. and Samuel R. Sommers (1997), "Inadmissible Testimony, Instructions to Disregard, and the Jury: Substantive Versus Procedural Considerations," *Personality and Social Psychology Bulletin*, 23 (10), 1046–1054.

Kemmelmeier, M. (2004), "Separating the Wheat from the Chaff: Does Discriminating Between Diagnostic and Nondiagnostic Information Eliminate the Dilution Effect?" *Journal of Behavioral Decision Making*, 17 (3), 231–43.

Koenker, Roger and Gilbert Bassett (1978), "Regression Quantiles," *Econometrica*, 46 (1), 33–50.

——— and Kevin F. Hallock (2001), "Quantile Regression," *Journal of Economic Perspectives*, 15 (4), 143–56.

Leemore, Dafny and David Dranove (2005), "Do Report Cards Tell Consumers Anything They Don't Already Know? The Case of Medicare HMOs," NBER Working Paper No. 11420, National Bureau of Economic Research.

Linn, S., S.C. Schoenbaum, R.R. Monson, B. Rosner, P.G. Stubblefield, K.J. Ryan (1982), "No Association Between Consumption and Adverse Outcomes of Pregnancy," *New England Journal of Medicine*, 306 (3), 141–45.

McFarland, Cathy, Adeline Cheam, and Roger Buehler (2007), "The Perseverance Effect in the Debriefing Paradigm: Replication and Extension," *Journal of Experimental Social Psychology*, 43 (2), 233–40.

Nisbett, R.E., H. Zukier, and R.E. Lemley (1981), "The Dilution Effect: Non-Diagnostic Information Weakens the Implications of Diagnostic Information," *Cognitive Psychology*, 13 (2), 248–77.

Pope, Devin G. (2009), "Reacting to Rankings: Evidence from 'America's Best Hospitals and Colleges,'" *Journal of Health Economics*, 28 (6), 1154–65.

Ross, Lee, Mark R. Lepper, and Michael Hubbard (1975), "Perseverance in Self-Perception and Social Perception: Biased Attributional Processes in Debriefing Paradigm," *Journal of Personality and Social Psychology*, 32 (5), 880–92.

Schul, Yaacov (1993), "When Warning Succeeds: The Effect of Warning on Success in Ignoring Invalid Information," *Journal of Experimental Social Psychology*, 29 (1), 42–62.

Slovic, P. and B. Fischhoff (1977), "On the Psychology of Experimental Surprises," *Journal of Experimental Psychology: Human Perception and Performance*, 3 (4), 544–51.

Steblay, Nancy, Harmon M. Hosch, Scott E. Culhane, and Adam McWethy (2006), "The Impact on Juror Verdicts of Judicial Instruction to Disregard Inadmissible Evidence: A Meta-Analysis," *Law and Human Behavior*, 30 (4), 469–92.

Tversky, Amos and Daniel Kahneman (1974), "Judgment Under Uncertainty: Heuristics and Biases," *Science*, 185 (4157), 1124–31.

Wooldridge, Jeffrey M. (2001), "Corner Solutions and Censored Regression Models," in *Econometric Analysis of Cross Section and Panel Data*, Chap. 16. Cambridge, MA: MIT Press, 517–49.

Zeithammer, Robert and Christopher Adams (2010), "The Sealed-Bid Abstraction in Online Auctions," *Marketing Science*, (published electronically August 11, 2010), [DOI: 10.1287/mksc.1100.0561].