The Welfare Impact of Consumer Reviews: A Case Study of the Hotel Industry†

Gregory Lewis
Microsoft Research and NBER

Georgios Zervas
Questrom School of Business
Boston University

July 16, 2016

Abstract

Platforms such as Yelp and TripAdvisor aggregate crowd-sourced information about users' experiences with products and services. We analyze their impact on the hotel industry using a panel of hotel prices, sales and reviews from five US states over a 10-year period from 2005-2014. Both hotel demand and prices are positively correlated with their average ratings on TripAdvisor, Expedia and Hotels.com, and such correlations have grown over our sample period from a statistical zero in the base year to a substantial level today: a hotel rated one star higher on all the platforms on average has 25% higher demand, and charges 9% more. A natural experiment in our data that caused abrupt changes in the ratings of some hotels but not others, suggests that these associations are causal. Building on this causal interpretation, we estimate a structural model of supply and demand with partially informed consumers, finding that in a counterfactual world without any review information, aggregate consumers surplus for potential travelers to these markets would fall by $124 million when prices are held fixed. Allowing prices to adjust leads to welfare conclusions that are sensitive to the modeling assumptions.

†We thank Michael Kurish for excellent research assistance. We also thank STR, and in particular Duane Vinson, for helping us with data acquisition.

1 Introduction

The internet has been one of the most significant innovations of recent times. It has made information about goods and services more readily available, allowing consumers to make more informed choices. Online review platforms are one important source of information. For example, TripAdvisor — the world's largest travel review platform — contains 320 million reviews. Yelp, which aggregates reviews for local businesses receives approximately 150 million visitors per month.[1] But the welfare implications and magnitudes of such information access are unclear. For example, when everyone in New York knows what the top ramen shops are, these businesses may raise prices or have long queues to ration demand, which in turn may actually leave consumers worse off.

We investigate how the increasing popularity of these information intermediaries has affected consumer choice, and how firms have responded to these changes in consumer behavior. The specific setting we investigate is that of the hotel industry. Given how the internet has transformed travel planning and booking, this an appealing setting to study.

Specifically, we ask three main questions. First, to what extent are information intermediaries helping consumers make better choices? Second, how are firms responding? Third, what are the welfare consequences of these forces combined? Our focus is on the market as a whole, rather than the effects of increased ratings on individual firms.

Our analysis is enabled by the combination of two novel data sources. The first data source we employ is a decade-long proprietary panel of hotel financial performance. The panel, which is at the month level, contains information on the revenue and demand (measured in room-nights) of roughly one half of all hotels that operated during our observation period in California, Oregon, Washington, Nevada, and Arizona (65% of all rooms available). Our second source of data is the entire historical record of online reviews for these hotels, which we collect from three major review platforms: TripAdvisor, Expedia (the world's largest online travel agent), and Hotels.com (an Expedia brand).

A unique feature of our data which allows us to identify the impact of information disclosure on consumer choice, is that it spans a period during which review platforms grew from insignificant sources of information to the major source of travel information. This allows us to analyze long-run trends. For example, one measure of the quality of a hotel is its average rating (across all three platforms) as of December 2014, when most hotels have a substantial number of reviews. Measured this way, the average quality of hotels that consumers stay

[1] See `http://www.tripadvisor.com/PressCenter-c4-Fact_Sheet.html` and `http://www.yelp.com/factsheet`.

at has been growing modestly over time, from around 3.8 to 4.0. This suggests that the impact of these platforms on consumer choice has been correspondingly modest: long before TripAdvisor and Expedia became popular information sources, consumers tended to choose hotels that are only of slightly quality (as measured by current ratings) than the ones they choose today.

Of course, there are other potential explanations for this modest growth: higher-quality hotels may have raised prices, so that they are chosen little more often than they were in the past. To analyze this more carefully, we estimate a logit demand system that controls for price, time-invariant hotel characteristics (through hotel fixed effects) and market-specific demand shocks (though market-year-month fixed effects). As a discrete choice model, it is also able to account for changes in competition arising from the information available to consumers about competing hotels. A maintained assumption of this model is that all consumers account for the review platform information when making choices. The reality is surely more complicated: more "tech-savvy" consumers and those with a strong taste for hotel quality are more likely to seek out this information. We think of our model as in some sense a "linearizaton" of a more complex mixture model.

We find that a a 1-point increase in a hotel's rating is associated with a 6.5% increase in demand and a 1.5% increase in price. But there is substantial heterogeneity in the effects across hotels and over time. Independent hotels are more sensitive than chains and franchises: the estimated demand increase is around 10% for independents compared to 5% for the other two groups. This is consistent with the idea that ratings are more influential for independent hotels because consumers do not have strong priors on their quality.

Luxury hotels also have bigger demand and price effects. This makes sense, as travelers to these hotels typically put a high premium on quality. We also find that the coefficients on ratings steadily increase over time in both the demand and price regressions, rising to 25% and 9% respectively by 2014.

Next, exploiting the long panel and the presence of multiple platforms, we show that the *difference* between the average TripAdvisor and Expedia rating is significantly correlated with demand, negatively towards the beginning of our sample period (when Expedia was significantly more well-known), and positively more recently (when TripAdvisor's growth in internet traffic has been much quicker than Expedia's). This points to a causal role for online reviews, as it would be surprising to find that the difference in reviews predicts market share if neither platform's rating has causal effects on demand.

As further evidence of causal relations, we exploit a natural experiment in our data. In June

2013, Expedia and Hotels.com merged their review collections, so that a hotel with different average ratings on the two sites prior to the merger displayed the same average rating after the merger. Because of the way ratings on these platforms are rounded to the nearest 0.1, it is often the case that a hotel's rating goes up by 0.1 on one platform, while remaining unchanged on the other, generating variation in the average displayed rating of the hotel. Treating the merger as exogenous, we estimate changes in hotel demand around around the date of the natural experiment using a difference-in-differences (DD) empirical strategy. Specifically, by comparing changes in demand for hotels whose ratings changed against hotels whose ratings were unaffected, we find that a 1-star increase is associated with a 26% increase in demand. These estimates are strikingly similar to our previous analyses.

The size of the treatment effect size we estimate may seem surprisingly large, but becomes more plausible when evaluated against the effort hotels would have to expend to achieve a 1-star increase and realize the associated gains in demand. For instance, consider a hotel with a 4-star average rating and 100 reviews, which resembles the median hotel in our data. Such a hotel would need an uninterrupted streak of two thousand 5-star ratings to increase its average rating by 1 star. Since most changes in our data are much smaller that 1-star, so are the changes they cause in hotel demand.

In the last part of the paper, we estimate a counterfactual in which the information available on the review platforms disappears. We assume that consumers form beliefs about hotel quality as they did in 2005, prior to these hotel review platforms becoming mainstream, using our demand system from that time period to infer those beliefs. We estimate how this affects choice, consumer surplus and revenue. We find that consumer surplus falls by $124 million in the counterfactual, holding prices fixed.

Next, we model price adjustment by hotels, using two different approaches. The first takes the same tack, using the 2005 pricing estimates to make inferences about how hotels would price in 2014 if there were no online reviews. The second takes a structural approach, inferring marginal costs from first order conditions, and then estimating Nash equilibrium prices under the status quo and counterfactual. We find that estimated equilibrium prices are similar in the two scenarios and close to existing prices, so that the change in consumer surplus is close to the earlier estimate (a fall of $107 million). By contrast, the predicted prices from the 2005 estimates are quite different to current prices, and when we calculate counterfactual consumer surplus using this approach, we find that it would *rise* by $546 million. Thus the results when accounting for price adjustment are sensitive to the modeling assumptions.

The impact of reviews and ratings on firms' sales has been a popular subject in the marketing and economics literatures. For example, closely related to our work, are a number of studies that estimate the impact of reviews and ratings on sales for various products and services including restaurants, books, and internet auctions (Chevalier and Mayzlin, 2006; Cabral and Hortacsu, 2010; Luca, 2011; Anderson and Magruder, 2012). These papers and several other papers have focused on estimating the treatment effect of a firm's online reputation on its financial performance, consistently arriving at the conclusion that reputation is a significant driver of sales. Our work differs from these papers in that we take a market-level view rather than a firm-centric one, accounting for how changes in information affect competition between firms within a market. In addition to estimating the impact of better ratings on the bottom line of firms, we are able to measure the welfare implications for consumers.

Along similar lines a number of papers measure the effect of quality disclosure on consumer choice. Jin and Leslie (2003) find that consumers respond to the disclosure of restaurant health ratings, and that restaurants respond by becoming cleaner; jointly these effects leads to a decrease in food-borne illness suggesting that health ratings had a positive impact on consumer welfare. Elfenbein et al. (2014) study eBay sellers and find that the extent to which a "top rated seller" helps attract more customers depends on how many other sellers who sell similar products also have the badge. Bai (2015) experimentally demonstrates that quality disclosure leads to increased prices and profits in market with high information asymmetries.

Finally, our work is directly related to the literature on measuring consumer surplus from the digital economy (Pantea and Martens, 2014). A difficulty in calculating the internet's consumer surplus is that much of what the internet provides is free to consumers. Therefore, it is difficult to measure surplus through the usual means of estimating demand elasticities for online services. To overcome this difficulty, Goolsbee et al. (2006) measure the consumer surplus of the internet by relating it to the opportunity cost of pursuing non-internet activities. A recent study by Brynjolfsson and Oh (2012) relies on similar methods. Both studies find significant gains associated with internet use, in the order thousands of dollars per consumer per year. In this paper, we take a different approach to measuring consumer surplus: while access to review platforms is free, consumers who consult them can make better choices for goods and services the purchase. By observing realized choices and comparing them against choice consumers would have made in the absence of review platforms, we can directly estimate the economic gains of online information intermediaries. Our work also contributes to our understanding of the specific channels through which internet use benefits consumers. The literature has primarily focused on benefits derived from decreased search costs and easier comparison shopping, which in turn have resulted in increased competition among firms, lower prices (Brown et al., 2002), and richer product assortments (Brynjolfsson et al., 2003). In this paper, we show that information sharing among consumers on crowd-sourced review platforms is another digital channel through which consumers have realized significant gains.

2 Theory

We begin the paper with a simple model of how information intermediaries such as TripAdvisor and Expedia can influence supply and demand in the hotel market. Consider a Hotelling model with two hotels A and B located at opposite ends of a unit interval. We assume linear transportation costs equal to t, so that a consumer located at $\theta \in [0, 1]$ earns utility $u_{A,t} = v_A - p_A - t\theta$ from staying at hotel A and $u_{B,t} = v_B - p_B - t(1 - \theta)$ from staying at hotel B. Hotels can be either high-quality ($v_j = H$) or low-quality ($v_j = L$), for $j = A, B$ and $H > L$. Consumers have a prior that high and low-quality are equally likely, and therefore in the absence of further information, have an expected gross utility (i.e. gross of price and transportation costs) of $1/2(H + L)$ for each hotel. To complete the model, we assume that consumers are uniformly distributed on $[0, 1]$ and the hotels have zero marginal costs.

Information intermediaries affect market outcomes by making hotel quality known to consumers through reviews and ratings. To understand the effects of intermediation, we consider two polar opposite scenarios: one in which consumers are uninformed about hotel quality (and therefore make decisions based on their prior), and one in which they are perfectly informed (corresponding to the case of perfect information intermediation). In the perfectly informed case, we will analyze a particular realization in which hotel A is high-quality ($v_A = H$) and hotel B is low-quality ($v_B = L$).

Figure 1 depicts the equilibria in these two different cases. We first consider the incomplete information case. The indifferent consumer is located at

$$\bar{\theta} = \frac{p_B - p_A + t}{2t}$$

and realized demand for the two hotels is $D_A = \bar{\theta}$ and $D_B = 1 - \bar{\theta}$. From the first order conditions of the firms' profit maximization problem it is easy to see that $p_A = p_B$ and hence $\bar{\theta} = 1/2$, *i.e.*, the firms split the market evenly as shown in the top panel of Figure 1.

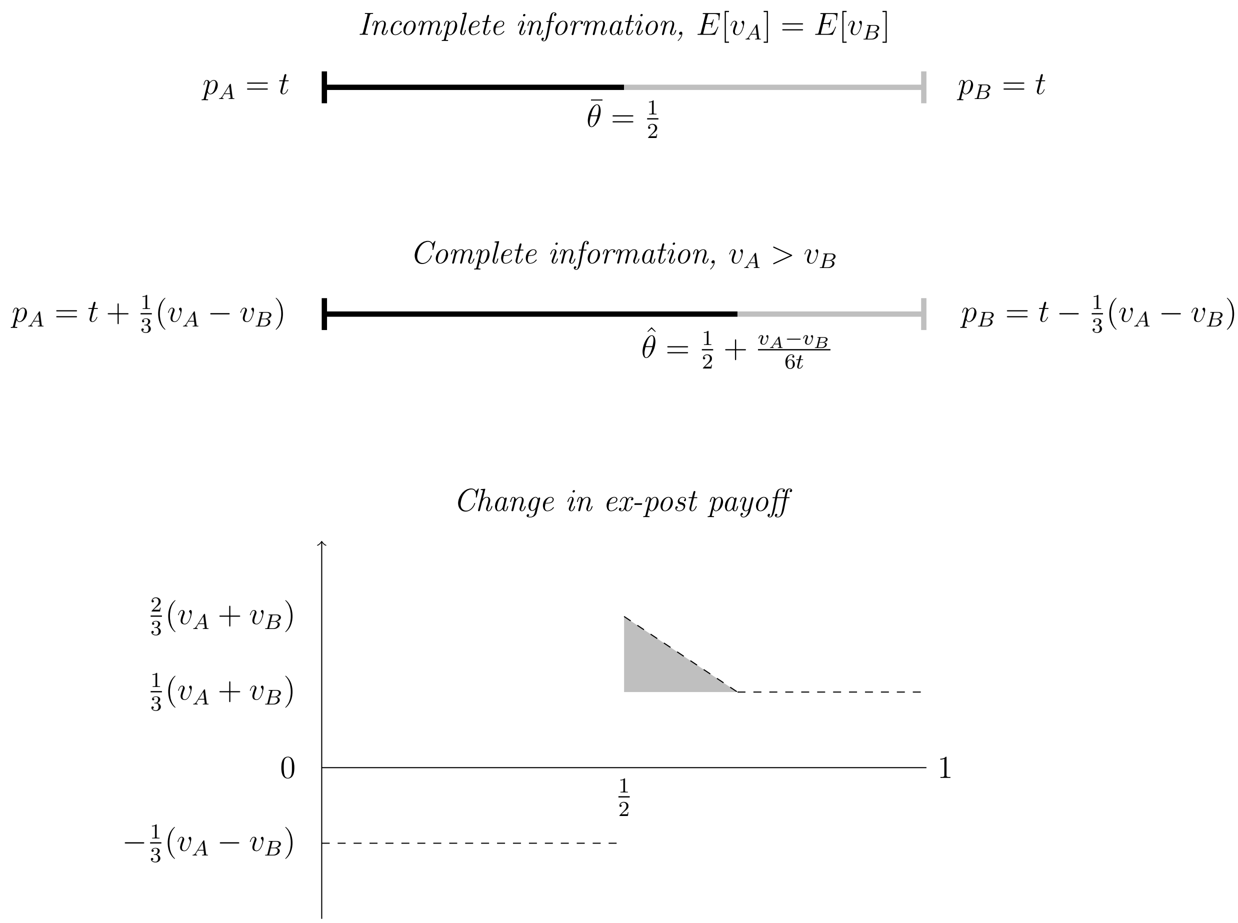

Figure 1: Welfare effects of information under a Hotelling model with consumers uniformly located on the unit interval, and linear transportation costs t. Firm A is high-quality, and firm B is low-quality. Firm A's market share is drawn in black, and firm B's in grey. Under incomplete information (top panel) consumers expect the firms to provide equal quality. Hence, the indifferent consumer is located at $\bar{\theta} = 1/2$ and firms set equal prices and split the market evenly. Under complete information (middle panel) consumers know that $v_A > v_B$. The indifferent consumer is now located at $\hat{\theta} > \bar{\theta}$. Firm A increases prices and captures a greater market share. The bottom panel displays changes in consumer welfare between the two information schemes as a function of consumer location.

We next consider the case where information intermediation permits the two hotels to differentiate by quality (middle panel of Figure 1). Consumers know that $v_A > v_B$ and the location of the indifferent consumer is now given by

$$\hat{\theta} = \frac{p_B - p_A + v_A - v_B + t}{2t}.$$

From the first order conditions of the firms' profit maximization problem it follows that optimal prices are $p_A = t + (v_A - v_B)/3$ and $p_B = t - (v_A - v_B)/3$. Under these prices we have that $\hat{\theta} = 1/2 + ((v_A - v_B)/6t) > \bar{\theta} = 1/2$. Therefore, in this example, information intermediation benefits higher-than-average quality firms both through higher prices and increased market shares (and conversely hurts lower-than-average quality firms).

For consumers, the effects of intermediation intermediation are more ambiguous. The bottom panel of Figure 1 plots changes in consumer welfare as a function of consumer location. First, we consider consumers located in $[0, 1/2)$, whose choice of hotel A is unaffected by information intermediation. Quality disclosure results in these consumers paying a higher price for the same hotel, while their transportation costs and the quality levels they experience remain the same. Therefore, information intermediation leaves consumers closer to firm A than firm B strictly worse off. Next, we consider consumers located in $(\hat{\theta}, 1]$, whose choice of hotel B also remains unaffected. Quality disclosure strictly benefits these consumers because it leads their chosen hotel to reduce prices. Due to the symmetric price adjustments of the two firms, the welfare gains of a consumer who continues choosing B are equal to the welfare losses of a consumer who continues choosing A, as shown by the dashed line in the bottom panel of Figure 1.

Last, we consider *switchers*, *i.e.*, consumers located in $[\bar{\theta}, \hat{\theta}]$ who switched from the low-quality to the high-quality hotel because of information intermediation. The benefits to switchers vary by location. Information intermediation produces the largest welfare gains for consumers who are equidistant to the two firms. The welfare difference of switching from hotel B to hotel A for a consumer located $\bar{\theta} = 1/2$ is $v_A - v_B$ while the difference in price is $1/3(v_A - v_B)$ for a net utility increase of $2/3(v_A - v_B)$. As the distance to the high-quality hotel increases the welfare gains of switching hotels decrease. The shaded triangle represents the aggregate increase in consumer welfare from information intermediation.

Notice that this welfare analysis is all ex-post (i.e for a particular realization of the quality levels of the two hotels). One could repeat this analysis for each of the remaining possible realizations (H, H), (L, H) and (L, L). The case (L, H) is identical to the present one up to labeling, and yields the same conclusion: not all consumers benefit from information intermediation, and it is the switchers who benefit most. When the hotels have equal ex-post quality, information intermediation has no real effects: prices are unchanged, consumers make the same choices, and welfare is unchanged. This conclusion is sensitive to the modeling assumptions: in the presence of an outside option (i.e. if the full market coverage assumption is dropped), total sales may rise in the (H, H) case and fall in the (L, L) case.

We can also compute the ex-ante welfare consequences of information intermediation (i.e. taking the average over all possible realizations). Firms gain $\frac{1}{2t}\left((v_A - v_B)/3\right)^2$, because with some probability they will have increased market power, and the returns to market power are convex in this model. Consumers who never switch (i.e. those outside of $[1 - \hat{\theta}, \hat{\theta}]$) receive the same expected consumer surplus, as the price changes cancel out when taking the expectation. The more elastic consumers in the center receive strictly higher consumer surplus, as when they optimally choose to switch from the the close low-quality hotel to the farther away high-quality hotel, they experience a net utility gain.

This simple static theory foreshadows our main results. As far as firms are concerned, the anticipated impact of information intermediation is clear: higher quality firms should increase prices and capture greater market shares. The welfare gains for consumers depend on how much they value quality relative to other characteristics (in this model, location). Those for whom quality is relatively unimportant will not change their decisions, and in expectation experience no change in consumer surplus. But consumers who value quality highly will benefit from information intermediation, even though firms with high-quality products raise prices. The aggregate change in consumer surplus thus depends on the distribution of taste for quality.

Our simple model of information intermediation yields a few conclusions that inform our empirical analysis. First, to measure the change in consumer surplus from information intermediation, we will need to measure the demand for hotels of different quality levels (whose empirical counterpart is average ratings). As a practical matter, the quality and attention paid to platforms like TripAdvisor may change over time, and so we will want to allow the demand for high-rated hotels to be time-varying. Given such estimates, we will need to infer the demand that we would obtain in a counterfactual world without information intermediation. This in turn will require making some assumption about the prior beliefs of consumers. Second, we will be unable to correctly measure the change in consumer surplus, revenues, and welfare due to information intermediaries without accounting for price responses. In this simple model, intermediation had no effect on average prices, but in general this will depend on how residual demand changes with perceived quality. Price responses can have substantial welfare impact, as they affect both switchers and non-switchers alike. We next turn to the data we will use to measure these objects.

3 Data

Our primary source of data for this study is a decade long monthly panel of hotel financial performance, which we obtained from Smith Travel Research (STR). The STR panel contains 5,944 hotels located in Arizona, California, Nevada, Oregon, and Washington. Approximately 45% of all hotels that operated in these five states during our observation period reported financial performance data to STR, and are thus included in our panel. The hotels in our panel are much more likely to be affiliated with a chain than to be independent: 90% of chains and 10% of independents report data to STR. For each hotel-month, we observe the number of room-nights available, the number of room-nights sold, and the total room revenue generated. Using these three variables, we also calculate average room prices and occupancy rates over time. In addition to these time-varying covariates, our data contains a rich attribute set covering both STR and non-STR hotels: hotel location at the ZIP code level, opening and closing dates (if any), price category (from Budget to Luxury), organizational form (chain, franchise, or independent), capacity, and the square footage of any meeting and business facilities. Our data masks the identities of individual hotels.

We augment the hotel financial performance data with a panel of consumer reviews from three major review platforms: TripAdvisor, Expedia, and Hotels.com. Figure 4 plots the (non-cumulative) number of reviews submitted to each the three platforms by year. In our data, we observe the entire history of 1-, 2-, 3-, 4-, and 5-star ratings for each hotel on each review platform. Our data does not contain the text of individual reviews to maintain hotel anonymity. In total, our reviews dataset contains 807,140 Expedia ratings, 1,410,488 Hotels.com ratings, and 1,544,883 TripAdvisor ratings. We aggregate ratings across platforms at the hotel-month level to match the financial performance data, defining the average rating $r_{j,t}$ of a hotel j in a specific month t as the sum of all individual ratings in reviews received by j up until time t across all three platforms, divided by the total number of reviews across the platforms.[2]

3.1 Descriptive evidence

We begin by offering some evidence on how consumers review hotels. Figure 2 shows the distribution of ratings on each review platform. In our data, negative reviews are rare – only 12% of ratings are below three stars. These distributions differ from some prior measurements, which found that extreme reviews are more prevalent than moderate ones, resulting in a J-shaped distribution (Hu et al., 2009). Reviewing patterns appear to be similar across the three platforms. Over time the correlation in the average reviews across

[2] In the appendix, we show that all our main results are robust to instead first defining platform-specific weights based on the platform share of reviews across all hotels up to time t, and then averaging average reviews on each platform according to these platform specific weights.

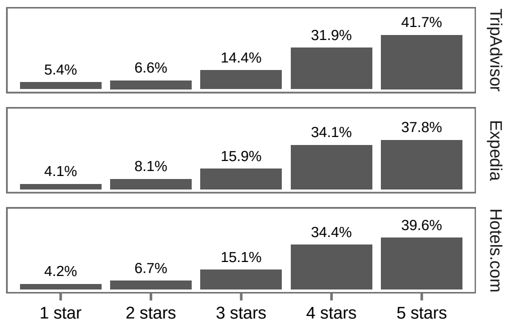

Figure 2: Distribution of ratings by review platform.

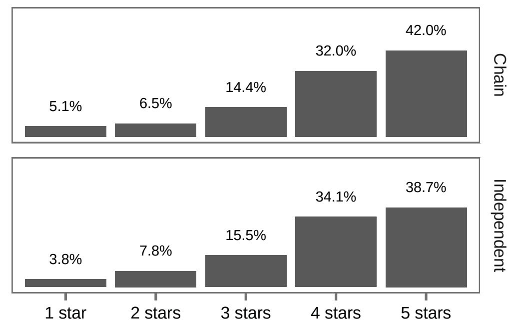

Figure 3: Distribution of ratings by hotel type.

platforms has also grown, so that the platforms largely "agree" about hotel quality.[3] The distribution of ratings is also similar across chain/franchise hotels and independents, despite the fact that independent hotels tend to be of average higher price ($122 for independents versus $80 for chain or franchise hotels).

Review platforms have been growing in popularity. Figure 4 shows the number of reviews submitted by year on each platform. Notice that the growth rate has increased in recent years, particularly on TripAdvisor. The prior literature has offered little evidence on the frequency with which consumers decide to leave a review. Figure 5 plots the annual probability of leaving a review on any of the three review platform in our data, which we define as total number of reviews over total number of reservations. Because we do not observe the latter quantity, we approximate it by dividing total room-nights by the average length of stay nights, which we assume to be two nights based on industry statistics.[4] We find that, while the rate at which travelers leave reviews has been growing, by 2014 about 2% of stays at independent hotels, and 1% of stays at chain hotels resulted in a review being left on Expedia, TripAdvisor, or Hotels.com. These figures are in sharp contrast to the ones seen on platforms like Airbnb, where 67% of guests leave a review on average (Fradkin et al., 2014).

While only a small fraction of consumers leave reviews, a much larger fraction consult them when making choices. Our interest in this paper centers in understanding how consumer choice has responded to the proliferation of online reviews. One way to measure whether consumers have gravitated towards high quality hotels is to form a measure of quality and then ask if the average quality of hotels stayed at has improved over time. We measure the

[3] Some of this growth in correlation is mechanical, since Expedia and Hotels.com merged their review databases in June 2013. See our discussion of the natural experiment in section 4.4 below.

[4] According to the American Hotel & Lodging Association, among leisure travelers staying at a hotel "50% spend one night, 26% spend two nights, and 24% spend three or more nights". Business trips are typically of shorter duration. These statistics suggest an average length of stay of approximately two nights. See `https://www.ahla.com/content.aspx?id=36332`.

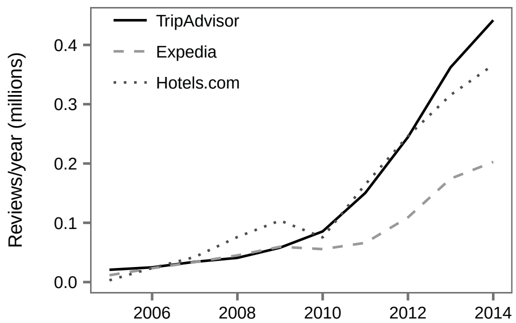

Figure 4: Number of reviews submitted by year.

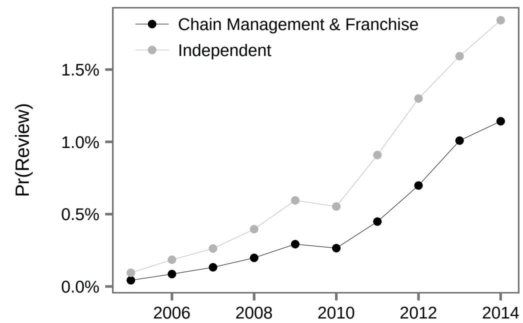

Figure 5: Probability of leaving a review on TripAdvisor, Expedia, or Hotels.com by year.

quality of a hotel using our best proxy for it, which is their average rating (over all review platforms) as of December 2014. We call this the hotel's terminal rating.

Figure 6 displays the year-by-year evolution of demand-weighted terminal ratings for chain and independent hotels. We observe that as review platform adoption grows, travelers appear to be making better choices, at least as measured by terminal hotel ratings. Interestingly, we observe a steeper trend for independent hotels. This observation coincides with our current understanding of organizational form and reviews. For instance, Luca (2011) shows that chains benefit less from reviews, and Mayzlin et al. (2014) show that chains are less likely to commit review fraud, in part because they benefit less from better reviews. Overall, chains accrue smaller informational gains from online reviews because they provide standardized services and benefit from brand recognition. Therefore, a user reading online reviews will on average acquire less new information about a chain compared to an independent hotel.

But there is no such thing as a free lunch. In Figure 7 we plot the corresponding changes in average terminal prices (i.e. Dec 2014 prices) paid by consumers. We use Dec 2014 prices rather than the prices for the relevant time periods to isolate the effect of consumers switching to high-rated hotels, and to avoid conflation with other sources of price changes. If higher rated hotels also tend to charge higher prices, we should see that these average terminal prices rise over time. This turns out to be true, but only for independent hotels. Again, organizational form may provide an explanation: independent hotels may be more easily able to adjust prices as consumers become aware of their quality, whereas chain hotels may be more constrained.

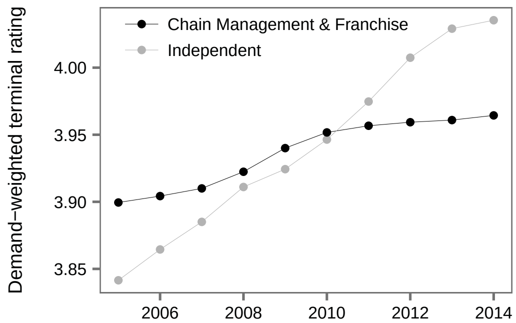

Figure 6: Demand-weighted terminal rating by year.

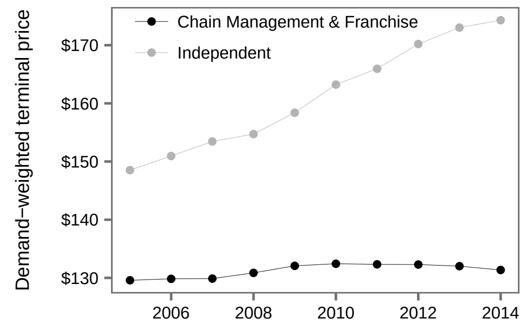

Figure 7: Demand-weighted terminal price by year.

4 Empirical strategy and results

We divide our econometric interrogation of the data into two parts: first, we examine how a hotel's ratings affect its demand, and second how its ratings affect the prices it charges. To analyze demand we will need a demand system with partially informed consumers. We offer a simple modification of the discrete-choice logit demand system in which consumers form beliefs about latent hotel quality using all the information at their disposal, and maximize expected utility. In the demand specification, in addition to the usual price endogeneity concerns, we need to address the potential endogeneity of ratings in order to plausibly identify a causal effect. For the most part, our strategy is simply to include flexible controls in order to limit such concerns. In section 4.3 we disaggregate our analysis and estimate platform and time specific coefficients (e.g. the coefficient on Tripadvisor ratings in January 2007). We show that these coefficients track engagement with the various platforms (as measured by Google Trends) closely, consistent with a causal interpretation of our earlier regressions. In section 4.4 we go further, exploiting a natural experiment based on a merger of the review databases of Expedia and Hotels.com to estimate a causal effect.

4.1 The impact of review platforms on hotel demand

We estimate a logit demand system with imperfectly informed consumers. Consumer i gets utility from choosing hotel j in market m at time t given by:

$$u_{i,j,t} = \gamma_{j,t} + \alpha p_{j,t} + \varepsilon_{i,j,t}$$

where $\gamma_{j,t}$ is the hotel quality, $p_{j,t}$ is price, and $\varepsilon_{i,j,t}$ is distributed iid extreme value type 1. However, consumers do not know the quality of the hotel. Instead they form beliefs based on the information in their information set $\Omega_{j,t} = (r_{j,t}, x_{j,t}, \xi_{j,t})$. $r_{j,t}$ is the hotel's average review (observed by the econometrician), $x_{j,t}$ are other observable hotel characteristics, and $\xi_{j,t}$ is a scalar latent unobservable. Notice that we do not allow price to affect their beliefs, as this creates well-known game theoretic complications that we feel poorly equipped to account for. We assume that the conditional expectation of quality is linear in these signals, so that:

$$E[\gamma_{j,t}|r_{j,t}, x_{j,t}, \xi_{j,t}] = \phi_t r_{j,t} + \beta_t x_{j,t} + \psi_t \xi_{j,t}$$

We allow for time-varying coefficients because the effect of the various signals on consumer beliefs may vary over time. This would be true, for example, if we specified a Bayesian model with Gaussian priors and signals, in which the conditional variance in average ratings given quality declined over time, so that consumers would optimally place higher weight ϕ_t on reviews over time.

Consumers maximize expected utility:

$$E[u_{i,j,t}|p_{j,t}, r_{j,t}, x_{j,t}, \xi_{j,t}, \varepsilon_{i,j,t}] = x_{j,t}\beta_t + \alpha p_{j,t} + \phi_t r_{j,t} + \psi_t \xi_{j,t} + \varepsilon_{i,j,t}$$

Then denoting the mean utility (before realization of the idiosyncratic error $\varepsilon_{i,j,t}$) as $\delta_{j,t}$, we get a linear specification:

$$\delta_{j,t} = x_{j,t}\beta_t + \alpha p_{j,t} + \phi_t r_{j,t} + \psi_t \xi_{j,t} \tag{1}$$

As is well known, in logit demand systems, the mean utility index $\delta_{j,t}$ is equal to $\log\left(\frac{s_{j,t}}{s_{0,m(j),t}}\right)$, where $s_{j,t}$ is the market-share of hotel j in time period t. The market-share is in turn defined as $\frac{q_{j,t}}{M_{m(j),t}}$ where $q_{j,t}$ is the number of room-nights hotel j sold in period t and $M_{m(j),t}$ is the number of potential buyers in market j at time t. We define $M_{m,t}$ as the maximum number of room nights ever available in that market (roughly equal to the supply of rooms at the market's peak). The term $s_{0,m(j),t} \equiv \left(1 - \frac{\sum_{j \in M} s_{j,t}}{M_{m(j),t}}\right)$ is the share of consumers choosing the "outside option" of non-purchase in the market $m(j)$ that hotel j is operating in. This gives us an estimating equation:

$$\delta_{j,t} = x_{j,t}\beta_t + \alpha p_{j,t} + \phi_t r_{j,t} + \psi_t \xi_{j,t} \tag{2}$$

where all the variables are observed except for the error term $\xi_{j,t}$.

In practice, we will amend this in various ways. First, we do not want our assumption on the market size $M_{m,t}$ to affect estimation. To do we will add market-year-month fixed effects $\tau_t \times m_j$ which flexibly control for unobserved demand shocks (such as seasonal demand variation and time trends) contemporaneously affecting all hotels within a specific market. Second, not all hotels are rated at all points in time, and so to deal with this we include a dummy $v_{j,t}$ indicating whether the hotel had been reviewed by period t, and add this as an additional control (with its own time varying coefficient). Third, in our baseline specification, we will assume that the only source of time series variation in beliefs is the ratings themselves, and control for all non-time-varying hotel characteristics with a fixed effect h_j (hotel observables are fixed in our data, so this is strictly more general than simply including a term $x_j\beta$).

This leads to the following specification, which we will refer to as the "fixed effects" demand specification:

$$\delta_{j,t} = \phi_{1,t}\left(r_{j,t} \times v_{j,t}\right) + \phi_{2,t} v_{j,t} + \alpha p_{j,t} + h_j + (\tau_t \times m_j) + \tilde{\xi}_{j,t}, \tag{3}$$

The error term $\tilde{\xi}_{j,t}$ is equal to the difference between $x_{j,t}\beta_t + \psi_t \xi_{j,t}$ and the fixed effect h_j. It has a structural interpretation as the time-varying part of beliefs about hotel quality that is not attributable to ratings.

There are two main endogeneity problems we should be concerned about here. The first is that ratings may be endogenous. For example, suppose that consumers only learn about hotel quality from word-of-mouth. Word-of-mouth is in the error term $\tilde{\xi}_{j,t}$. If ratings are correlated with word of mouth, we will estimate positive coefficients $\{\phi_t\}$ even though they have no causal role in determining demand. This is central to our paper, and so we will spend quite a lot of time exploring the issue of ratings endogeneity below.

A second and more standard problem is that of price endogeneity. This is a more challenging problem than usual due to modern revenue management techniques: demand shocks are quickly transmitted to prices because hotels increase the price of their rooms as they sell out. Our approach is to estimate α separately using a strategy based on supply-side moment conditions, and then plug our estimate into (3) and proceed. As the actual procedure is somewhat involved, we summarize it quickly below and refer interested readers to the appendix for more details. Our estimates of the \$ value consumers place on higher rated hotels are not particularly sensitive to our estimate of α within a reasonable range around our estimate.[5]

[5] We have re-estimated the demand system for a range of α's corresponding to average residual demand elasticities between -1.4 and -2.2 and obtained estimates that vary from \$14 per star to \$17 per rating point — our preferred estimate of α corresponds to an average elasticity of -1.55.

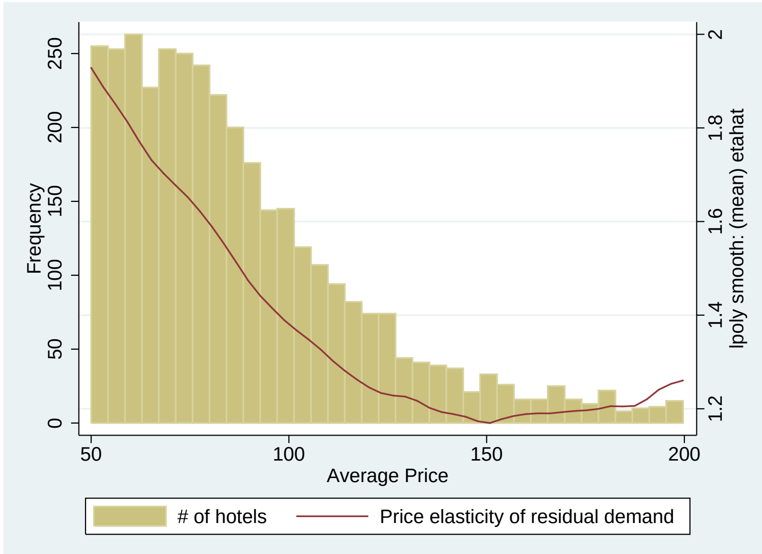

Figure 8: Price elasticity estimates.

Price endogeneity. The easiest approach to dealing with endogenous prices is to instrument for them. Unfortunately, good price instruments are hard to come by in the hotel market. Most variation in prices is driven by correlated demand shocks, implying that Hausman-style instruments are inappropriate. Since we allow for market-year-month fixed effects, valid cost instruments would have to shift individual hotel marginal costs within a market rather than marginal costs as a whole, and such instruments are difficult to find. BLP instruments based on the average characteristics of competitors do not vary much in our data, especially at the market-year-month level. We are not the first to encounter this problem. Koulayev (2014) also seeks hotel price instruments but settles for a rich set of controls instead, as he is met with the same problems we discover here.

Instead, we estimate the price coefficient α by enforcing a set of supply-side moments that require that marginal revenue, assessed at the prevailing prices, should be equal across the high season and low season (where high season and low season are defined at the state level). This condition is implied by price optimization by the hotel (such optimization is plausible given the aforementioned revenue management) and an assumption that marginal costs are constant over the sample period. Using a flexible semi-parametric demand specification, we recover estimates of the price elasticity of residual demand at each point in a grid of average prices and quantities.

Figure 8 illustrates the results of this procedure. The histogram shows the frequency of hotels whose average price over the sample period falls into each range. The distribution is clearly skewed, with the median hotel charging around $100 but a tail of hotels charging much higher prices. Superimposed on this is a local polynomial plot of the estimated price elasticity of residual demand. The estimated residual demand elasticity tracks the frequency of hotels quite well. This makes sense since residual demand will depend in part on the number of close competitors, and a reasonable notion of closeness is given by competitors who charge similar prices and are of similar size. Moreover, the absolute levels of the estimated elasticities are reasonable, ranging from -1.2 to -2 (the y-axis plots the absolute value).

In a logit demand system, the elasticities take the form $\eta_{j,t} = \alpha p_{j,t}(1 - s_{j,t})$. We estimate α by moment matching, as $\frac{\overline{\eta_{j,t}}}{\overline{p_{j,t}(1-s_{j,t})}}$, where $\overline{\eta_{j,t}}$ are the sample average estimated elasticity and $\overline{p_{j,t}(1 - s_{j,t})}$ is the sample average of that quantity. This implies an estimate $\hat{\alpha} = -0.0155$, implying an average elasticity of around -1.55 (since average prices are $\approx$ \$100 and shares are small). Further details are given in the appendix.

Fixed effects demand specification. Given our estimate $\hat{\alpha}$, we re-arrange Equation 3 as:

$$Q_{j,t} = \phi_{1,t} r_{j,t} \times v_{j,t} + \phi_{2,t} v_{j,t} + h_j + \tau_t \times m_j + \tilde{\xi}_{j,t}, \tag{4}$$

where $Q_{j,t} \equiv \delta_{j,t} - \hat{\alpha} p_{j,t}$. We will refer to $Q_{j,t}$ as the *adjusted log quantity*. The coefficient of interest is ϕ_1, which in the absence of ratings endogeneity can be interpreted as the impact of review platforms on demand.[6]

We present our initial results in the first column of Table 1. All our specifications double cluster errors at the market-year and hotel level to account for correlation in hotel prices and demand. In the first column, we show results assume that the ϕ coefficients are not time-varying. We find a significant effect of review platforms on hotel demand: a one-star increase in a hotel's cross-platform rating is associated with a 6% increase in room-nights sold. The effect size we estimate is in line with previously reported estimates (Luca, 2011; Anderson and Magruder, 2012). This estimate captures the average effect of review platforms on hotel demand during our 10-year observation period.

We next investigate heterogeneity of effects. Column (2) repeats the specification but interacts ratings with hotel organizational form (e.g. franchise). Chain and franchise hotels have an informational advantage over independent hotels derived from brand recognition, standardized service, and little variation in quality from location to location. Lacking these advantages, one would expect that high-quality independent hotels should benefit more from

[6]Estimates of ϕ_2 are omitted due to space constraints, but are well approximated by $\hat{\phi}_2 = 2\hat{\phi}_1$, as though consumers infer a rating of 2 for a hotel with no rating. The constant does decline slightly over time; being unrated in 2014 is a worse signal than being unrated in 2005.

an additional information channel for consumers to learn about product quality. We find that the impact of review platforms is substantially larger for independent hotels than chains – 10.4% vs 4.3% for a 1-star increase in rating. This analysis suggests that review platforms have helped narrow the informational gap between chains and independents.

In column (3) we interact with price class (where the order of class ranges from economy through midscale, upper midscale, upscale, upper upscale to luxury). We would expect that reviews matter more at the top, since consumers who can afford to pay high prices are also more discriminating. This is indeed what we find: the estimated coefficients have the expected order, and ratings matter most for luxury hotels.

Our final specification relaxes the assumption that the coefficients are stable over time. As noted earlier, we would expect increasing coefficients, because the informational content of review platforms has increased over time as more reviews have come in. For instance, in 2005 66% of the hotels in our data had at least one TripAdvisor review. By 2014 this number had increased to 99% of hotels. There has also been increased engagement with these platforms (see section 4.3 below for evidence on this). In column (4) we interact $r_{j,t} \times v_{j,t}$ and $v_{j,t}$ with year dummies. As expected, we find that review platforms have become more influential over time: by 2014 the impact of a 1-star increase on sales is 25%. The pattern is not entirely monotone: there is a slight dip in 2009-2010, corresponding to the recession in those years. As we'll see below, this may be an artifact of the fixed effects specification.

An alternative demand specification. In the fixed effects specification, we encode all of the consumer's non-time-varying beliefs in a fixed effect. One disadvantage of this specification is that it doesn't allow us to see if consumers are substituting towards online reviews and away from other information sources. So instead we now explore an alternative specification that imposes different restrictions on (2).

Let us assume that in 2005, consumers placed no weight on online reviews (i.e. $\phi_{1,2005} = \phi_{2,2005} = 0$). This assumption seems reasonable given the fixed effects results. Let us also drop the observables $x_{j,t}$ from the model, and write the scalar unobservable for 2005 as a mean plus a deviation:

$$\xi_{j,t} = \xi_j + e_{j,t} \quad , \quad t \in 2005$$

Finally, assume that $\psi_t = 1$ in 2005 (the scaling is without loss of generality, since the variable is latent — the restriction is that the coefficient is constant that year). Then the mean utility index in 2005 is just equal to $\delta_{j,t} = \alpha p_{j,t} + \xi_j + e_{j,t}$, and we can estimate ξ_j as the hotel-specific sample mean of $\delta_{j,t} - \hat{\alpha} p_{j,t}$.

This estimate $\hat{\xi}_j$ gives us a measure of the perceived quality of hotel j *before* online reviews were reliable and influential. This measure will prove useful in the counterfactual, as it tells us what consumers knew prior to the advent of online reviews.

For now, we will simply treat this as a part of the consumer's information set, estimating:

$$Q_{j,t} = \phi_{1,t} r_{j,t} \times v_{j,t} + \phi_{2,t} v_{j,t} + \psi_t \hat{\xi}_j + \tau_t \times m_j + \tilde{e}_{j,t}, \tag{5}$$

where $e_{j,t} = \psi_t(\xi_{j,t} - \hat{\xi}_j)$. Relative to the fixed effects specification, this is less flexible in the sense that it uses the estimate of $\hat{\xi}_j$ to control for all omitted hotel characteristics; but also more flexible in that it allows dependence on $\hat{\xi}_j$ to be time-varying.

The results are presented in Table 2, repeating all the specifications of Table 1. The results are extremely similar, suggesting that the additional flexibility of the fixed effect model is unnecessary. In fact the more parsimonious model delivers more sensible results in some respects: the ratings coefficients are monotone in year, as one would expect. This estimation strategy requires a balanced panel (since ξ_j is only defined if the hotel is present since 2005), and so the difference may be due to the experience of new entrants during the recession years (who are excluded from the present regression).

Notice also that the coefficients on ξ_j are statistically indistinguishable from one in every year. This is very informative: if consumers had access to information sources in 2005 that were time-varying and autocorrelated, then one would expect the coefficient on ξ_j to decay over time. The fact that it doesn't suggests that the latent signal can be decomposed into an almost perfectly persistent term (e.g. hotel characteristics observed by consumers) and possibly an additional iid noise term (e.g. short-term marketing efforts). Ratings endogeneity would thus arise only from correlation between this noise term and ratings, rather than from a more persistent source.

4.2 The impact of review platforms on hotel prices

As our simple theory of Section 2 argues, when hotel quality is disclosed, higher quality hotels should respond by raising prices as long as they have some market power. Similarly, lower quality hotels should respond by lowering prices. This response allows hotels to capture some of the welfare that review platforms generate.

Our next set of regressions investigates this hypothesis. The specification we estimate takes the following form:

$$p_{j,t} = \phi_1 r_{j,t} \times v_{j,t} + \phi_2 v_{j,t} + h_j + \tau_t \times m_j + \varepsilon_{j,t}, \tag{6}$$

which is similar to Equation 5, except that now the dependent variable we analyze is hotel prices. Unlike our demand specifications, these regressions are "reduced-form" in the sense that they are not derived from a model of optimizing behavior by firms.

We present our results in the first column of Table 3. We find that a 1-star increase in a hotel's cross-platform rating is associated with a \$1.5 increase in price for a room-night. Given that the average cost of room night in our data is about \$100, this represents roughly a 1.5% increase. This result suggests that hotels are capturing some of the welfare generated by review platforms. Even though more consumers are on average making better choices, these choices tend to cost more than they would in the absence of review platforms. The effect is relatively modest though.

In subsequent columns we again investigate heterogeneity, finding bigger effects for independents (\$3.5 or 2.4%). Again, this result is consistent with independent hotels being at an larger informational disadvantage prior to the adoption of review platforms. We also find that higher-class hotels charge more when they have good reviews (e.g. for luxury hotels, \$11 or 8.5%). This suggests that high-priced hotels are best positioned to exploit a reputation for high quality, which makes sense.

We also estimate a model that includes interaction with year dummies to measure changes in price responses over time. We present these results in column (4) of Table 3. We find a pattern of increasing price responses, consistent with our earlier observation that the demand response to reviews is increasing over time. By the end of the sample period the estimated price response is more substantial (\$9 or 9%).

4.3 Evidence on causality from trends in platform engagement

We would like to argue that the estimated effects of reviews on demand and prices are causal. As we have argued earlier, based on the prior regressions and comparison of the magnitudes to prior work in this literature, this seems plausible. Still, any time-varying omitted demand shock that is correlated with hotel reviews will bias the review coefficients.

A simple story in favor of such bias might be the following: hotel quality is unknown by consumers, but highly correlated with average reviews (*i.e.*, crowdsourced reviews are generally a pretty good quality measure). Consumers learn about hotel quality through many

Figure 9: Coefficients on ratings by platform and ratings difference across platform.

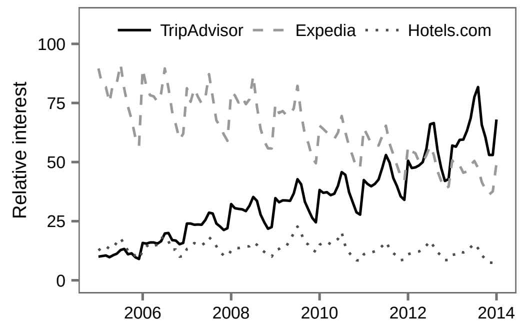

Figure 10: Google Trends by review platform.

sources: platforms such as TripAdvisor and Expedia, competing platforms not in our data (*e.g.*, hotel ratings offered by Google), their own research conducted using online and offline media, and marketing efforts, which may be disproportionately carried out by high quality hotels. Most of these information sources are omitted in our regressions, and hence appear in the error term, and yet are correlated with reviews through hotel quality.

We believe that online review platforms are an important or perhaps even the main information source for consumers, and so although the results are biased upward, the bias may not be too large. One way to assess this is to make use of the multiple platforms in our data. To do so, we run a disaggregated version of the demand system in Equation 5, allowing the average rating on TripAdvisor $r^T_{j,t}$ to enter separately from the average rating on Expedia/Hotels.com $r^E_{j,t}$, interacting each variable with year.[7] The difference $\Delta r_{j,t} = r^T_{j,t} - r^E_{j,t}$ is on average positive in the early period of our data (*i.e.*, TripAdvisor reviews are more favorable), but is zero on average in later years. Still, there is substantial variance in ratings across platforms, with the standard deviation of $\Delta r_{j,t}$ equal to 1.75.

The estimated coefficients on $r^T_{j,t}$ and $r^E_{j,t}$ by year are plotted in Figure 9. The coefficients on TripAdvisor increase substantially over time, from a baseline of almost zero in 2005 to 0.18 by 2014. By contrast, the coefficients on Expedia display much more gradual growth, moving from 0.04 in 2006 to 0.1 by 2014. Figure 10 shows the corresponding Google Trends data for search terms related to TripAdvisor, Expedia and Hotels.com. Notice how TripAdvisor's increasing share of searches is correlated with the increase in coefficients on $r^T_{j,t}$, which is suggestive of a causal role; while the gradual decline in Expedia searches is matched by the gradual increase in the coefficients on $r^E_{j,t}$.

Based on this, we run a slightly different specification in which we omit both $r^T_{j,t}$ and $r^E_{j,t}$, and

[7] We treat Expedia and Hotels.com as a single platform here, since they merged during our sample period (see section 4.4 below), and started showing consumers average ratings based on their joint stock of reviews.

instead include $\Delta r_{j,t}$ (interacted with year) as a covariate in (5). Notice that if the review platforms had no causal role — *i.e.*, average reviews are merely noisy versions of a quality measure that consumers learn from other information sources — one would expect a zero coefficient on $\Delta r_{j,t}$.[8] Instead, what we find is shown by the dotted "Difference" line in (9): a growing coefficient on the difference between the average Tripadvisor and Expedia review, negative in the first few years (when Expedia had a greater share of searches), and rising to 0.1 in the last year of our data. The difference between the two is statistically significant at 1% over the period 2011-2014. If one were to assume that the average Expedia review was a good proxy for all other information sources, one would conclude that an additional rating point on TripAdvisor, by itself, would increase demand in 2014 by 10%. This does suggest some causal effect of these review platforms, independent of other information sources.

4.4 Evidence on causality from a natural experiment

To further reinforce our causal claims we exploit a natural experiment occurring in our data. In March 2012, Expedia announced its intention to to incorporate reviews from its sister company Hotels.com in its review collection.[9] While Expedia announced these plans in March 2012, the merger did not take place immediately. Unfortunately, we could not find any public announcement from Expedia disclosing the date at which the change was implemented. Instead, we turned to the Internet Archive (IA). The Internet Archive provides access to historic snapshots of billions of web pages going back decades in time. By browsing through IA snapshots of Expedia web pages, we discovered that Expedia started showing Hotels.com reviews in June 2013. We arrived at this conclusion based on two markers: first, during June 2013 Expedia started displaying reviews marked with the phrase "Posted [DATE] on Hotels" (where "[DATE]" is, *e.g.*, "May 30, 2014"); second, between May and June 2013 the reported total review count for each hotel increased abruptly.

The merger affected the ratings displayed on Expedia and Hotels.com. Since Expedia and Hotels.com round their ratings to the nearest 0.1, in many cases where the two initial ratings were similar, the merger had the effect of raising or lowering the rating on one platform by multiples of 0.1, having no effect on the other. For instance, suppose that prior to the change the set of ratings Hotels.com is $H = \{4, 4, 4, 4, 5\}$ and the set of ratings on Expedia is

[8] Formally, suppose $r^p_{j,t} = \gamma_{j,t} + \varepsilon^p_{j,t}$ for $p = T, E$, and that consumers know $\gamma_{j,t}$ from other sources. Then $\Delta r_{j,t} = \varepsilon^T_{j,t} - \varepsilon^E_{j,t}$, and regardless of the distribution of $\varepsilon^p_{j,t}$, there is no information for consumers in $\Delta r_{j,t}$ and hence it should play no role in determining demand.

[9] See `https://viewfinder.expedia.com/news/expedia-overhauls-hotel-reviews-consumers-can-now-sort-verified-reviews-by-shared-interest/`

$E = \{4\}$ resulting in rounded ratings of 4.2 and 4 respectively. Post merger (and assuming no new reviews) the rounded rating on Hotels.com remains 4.2 stars, while the Expedia rating increases from 4 to 4.2 stars.

It is this kind of plausibly exogenous variation we exploit in the analysis below. Notice that the rounding is a necessary component for this merger to generate useful variation: the average rating across all three platform, $r_{j,t}$, is by construction unchanged by the merger (the underlying set of reviews remains the same), but the average rounded rating $\tilde{r}_{j,t}$ isn't.

Using our prior notation, let $r^p_{j,t}$ for $p \in \{E, H, T, EH\}$ be a hotel j's average rating on platform p at time t, and $n^p_{j,t}$ its (cumulative) number of reviews. We will use EH to represent Expedia/Hotels.com post-merger. Additionally, let $\tilde{r}^p_{j,t}$ denote hotel j's *rounded* average rating.[10] Next, define a hotel's cross-platform volume-weighted average rounded rating as:

$$\tilde{r}^P_{j,t} = \frac{\sum_{p \in P} n^p_{j,t} \tilde{r}^p_{j,t}}{\sum_{p \in P} n^p_{j,t}}. \tag{7}$$

Let $P_1 = \{E, H, T\}$ represent the setting where the review platforms operate independently, and $P_2 = \{EH, T\}$ represent the post-merger setting. Define the difference:

$$\Delta \tilde{r}_{j,t} = \tilde{r}^{P_2}_{j,t} - \tilde{r}^{P_1}_{j,t}, \tag{8}$$

which can be interpreted as the difference between ratings rounded following the merger and the counterfactual rounded ratings had the merger not taken place. Then, our estimating equation is:

$$Q_{j,t} = \phi_1 \Delta \tilde{r}_{j,t} \times \text{Post}_t \times v_{j,t} + \phi_2 v_{j,t} + h_j + \tau_t \times m_j + \tilde{\xi}_{j,t}, \tag{9}$$

where Post_t is a binary indicator for post-merger time periods. This can be thought of as a difference-in-differences (DD) specification where the "treatment" (a change in average rounded ratings) is continuous and generated by this merger. To consistently estimate a causal effect, we need that treatment assignment is uncorrelated with the error term. This is plausible given that the treatment is generated by a supply-side change (a merger).

Table 4 shows the characteristics of hotels whose average ratings were affected by the merger (the "rating shift" column), as compared to those who experienced no change. The treatment group, which is larger, comprises hotels who experienced any kind of change, which includes as particular cases having their rating rise on Expedia and fall on Hotels.com, or rise on Expedia and stay constant on Hotels.com (strictly rising or falling on both is mathematically

[10]Recall that Expedia and Hotels.com round ratings to 0.1-star increments whereas TripAdvisor rounds to half-star increments.

impossible). We see that the treatment and control groups are balanced on the observables, so that one might reasonably interpret any effect we find as an average treatment effect, rather than a local average treatment effect.

We estimate this regression using data from a 6-month window on either side of the merger. Results are presented in Table 5. Columns (1) and (3) show the results for adjusted log quantity and prices. Columns (2) and (4) report outcomes after adding the counterfactual rating as an additional control. This is a robustness check in case the treatment assignment is correlated with hotel's ratings. In all four cases we find an effect whose magnitude is similar to (and statistically indistinguishable from) the corresponding OLS estimate for 2013, as reported in the fourth column of tables 1 and 3.

As a final robustness check, we run a Placebo test, where we repeat this exercise for every other month in 2013. The results are reported in table 6. Consistent with our understanding of the merger's timing, we find a statistically significant effect on adjusted log quantity only in June.

5 Counterfactuals

Our empirical work suggests that both demand and prices are significantly affected by online ratings, especially in more recent years. These associations are likely causal.

We would now like to estimate the welfare impact of these review platforms by simulating a counterfactual world in which consumers do not have access to review data, and compare it to the status quo as of 2014 (when review platforms appear to be most valued by consumers).

The conceptual framework was outlined earlier in Section 2. In a world without review data, consumers will form different beliefs about quality and make different choices. We already saw this in the time series in Figure 6, where we showed that consumers gradually selected hotels with higher rating (measured as of Dec 2014) over time.

To quantify these beliefs, we use the demand system estimates from 2005, which was largely a world in which review platforms played no role. Specifically, we will take our second demand specification, in which consumers know a latent and fixed signal ξ_j for each hotel, and assume that they behave as they did in 2005, when that was the only signal available.

Given these beliefs, and holding prices fixed, we can calculate consumer choices, hotel revenue and consumer surplus. But as our model suggests, we should expect hotels to adjust prices. We offer two different approaches for doing this.

The first is to once again use the 2005 data. In 2005, firms set prices for a world without online reviews. We assume that in the counterfactual world they would employ the same strategies as they did in 2005, and model their pricing using a slightly more flexible version of the regression in the fourth column of table 3. The advantage of this approach is that it is straightforward and plausible: we see what firms would do in the counterfactual based on what they did in a time period similar to this one. The disadvantage is that the market structure has changed over time, and so the similarity may not be enough to make this comparison plausible.

Our second approach is to follow the standard empirical industrial organization playbook, and compute complete information Nash equilibrium prices, both with the current data and in the counterfactual (for an apples-to-apples comparison). To do this, we must first estimate marginal costs. We describe how we do this below.

There is an important data issue to address first though. We need estimates of market shares and prices for all hotels in our target markets for the year 2014. But since our STR dataset only covers 50% of hotels, we do not have all the required data. Ignoring the remaining hotels will dramatically oversimplify the market structure. So to deal with this, we "complete" the dataset consisting of both STR and non-STR hotels (which has missing ratings and revenue data for the non-STR hotels) by imputing a value for any necessary missing variable. In the imputation procedure, we group markets together by their fixed effect in the price regression into four groups, and then match hotels based on market-group, age, price class, size, amenities and organizational form (i.e. chain, franchise or independent). Similarly, hotels that are not present in 2005 do not have an estimated ξ_j. For those hotels we also need to impute a ξ_j. If hotels with missing data are systematically different from those with complete data conditional on these characteristics, we will introduce systematic error into the counterfactuals.

Simulations without price responses. Recall that the quality of a hotel is $\gamma_{j,t}$. In an abuse of notation, let us also denote by $\gamma_{j,t}$ our best estimate of the consumer's perceived quality of hotel j in month t (for months in 2014). We obtain these from the estimates from the fourth column of table 3. We also compute *counterfactual* perceived qualities for each hotel $\gamma^c_{j,t}$ by first predicting perceived quality using the 2005 coefficients instead. Since the counterfactual fixed effect for each market-month in 2014 is not identified by this strategy, we assume that the perceived quality in each market is unchanged, adjusting our initial estimates of $\gamma^c_{j,t}$ by a market-month specific constant so that the average value of $\gamma_{j,t}$ and $\gamma^c_{j,t}$ are equal in each market-month. This adjustment is implied by rational expectations:

removing a signal from consumer's information sets should not change their mean beliefs.

Given these estimates, we immediately get mean expected utility $\delta_{j,t} = \gamma_{j,t} + \alpha p_{j,t}$ and counterfactual mean expected utility $\delta^c_{j,t} = \gamma^c_{j,t} + \alpha p_{j,t}$. Status-quo and counterfactual market shares immediately follow from the logit formula. Revenues are calculated as shares $s_{j,t}$ and $s^c_{j,t}$ times observed prices $p_{j,t}$ times market size $M_{m(j),t}$.

Under the status quo, expected consumer surplus for an individual i in market m is given analytically by:

$$E[CS_{m,t}] = \frac{1}{\alpha}\log(1 + \sum_{j \in M} \exp(\delta_{j,t})) + C$$

where C is an unidentified constant that will be differenced out when comparing to the counterfactual.

In the counterfactual, consumers make choices according to $\delta^c_{j,t}$, but actually experience mean utility equal to $\delta_{j,t}$. We can still derive an analytic formula for this case. Let $j^*_i = \arg\max_j \left(\delta^c_{j,t} + \varepsilon_{i,j}\right)$. Then we have:

$$\begin{aligned} E[CS^c_{m,t}] &= E[\delta_{j^*_i,t} + \varepsilon_{i,j^*_i}] \\ &= E[\delta^c_{j^*_i,t} + \varepsilon_{i,j^*_i}] + E[\delta_{j^*_i,t} - \delta^c_{j^*_i,t}] \\ &= \frac{1}{\alpha}\log(1 + \sum_{j \in M} \exp(\delta^c_{j,t})) + C + \sum_j s^c_{j,t}(\delta_{j,t} - \delta^c_{j,t}) \end{aligned}$$

where the second line just adds and subtracts the true utility $\delta_{j,t}$, and the third line once again uses the consumer surplus formula.

Simulations with reduced form price responses. As outlined earlier, one approach for inferring counterfactual prices is simply to assume they are set in a similar way to how they were in 2005. There are many ways to do this. We choose to do something analogous to what we did for mean utilities, using the 2005 coefficients in the fixed effect regressions of Table 3. In practice we estimate a more flexible model than we have space to report, interacting year with ratings with class, so that the difference in prices between the status quo and counterfactual for a hotel with a given rating may vary with class of hotel. Again the question of how to set the counterfactual fixed effect arises, and again we choose to adjust the counterfactual prices so that the mean prices in each market under the status quo and counterfactual are equal. In contrast to before, there is no principled reason for making this particular assumption, other than that it is clear and seems reasonable. Given these counterfactual prices, we proceed as in the section above.

Simulations with equilibrium prices. To estimate equilibrium prices, both under the status quo and counterfactual, we need to fill in a model of the supply side. We assume that marginal costs are constant in 2014 up to an idiosyncratic error:

$$c_{j,t} = \bar{c}_j + \omega_{j,t}$$

Under logit demand, the supply side first order condition can be written as:

$$s_{j,t} + \alpha s_{j,t}(1 - s_{j,t})(p_{j,t} - c_{j,t}) = 0$$

Because the error is idiosyncratic we have:

$$E[s_{j,t} + \alpha s_{j,t}(1 - s_{j,t})(p_{j,t} - \bar{c}_j)] = 0$$

and replacing the LHS of this expression with the empirical averages for 2014 yields the following estimate of $\bar{c}_j$:

$$\bar{c}_j = \frac{\alpha \sum_t p_{j,t} s_{j,t}(1 - s_{j,t}) + \sum_t s_{j,t}}{\alpha \sum_t s_{j,t}(1 - s_{j,t})} \tag{10}$$

Given these costs, and the perceived status quo and counterfactual qualities $\gamma_{j,t}$ and $\gamma^c_{j,t}$, we calculate the complete information Nash equilibrium prices in each case for each market-year-month by finding a solution to all the first order conditions given the estimated costs.

The status-quo equilibrium prices are not equal on average to those we observe in the actual data, which would suggest some failure of the assumptions (e.g. demand is not logit, firms are not optimizing, firms have different information sets), though they are close (the difference in means is \$1). Given this, we proceed by comparing the consumer surplus, revenues and prices under the *simulated* prices and status quo to the *simulated* counterfactual objects and taking differences.

Counterfactual results. Consider first figure 11, which shows some of the inputs to the counterfactual exercises. The shaded dots are the difference between counterfactual and status quo perceived quality for hotels of various ratings levels (ratings are rounded to a half-star, the category 2 includes hotels that are unrated and have ratings less than 2). Counterfactual utility rises for poor quality hotels, and falls for high quality hotels, linearly (this is a consequence of the linear empirical specification).

The unshaded dots show the price change using the reduced form estimates. Here we see

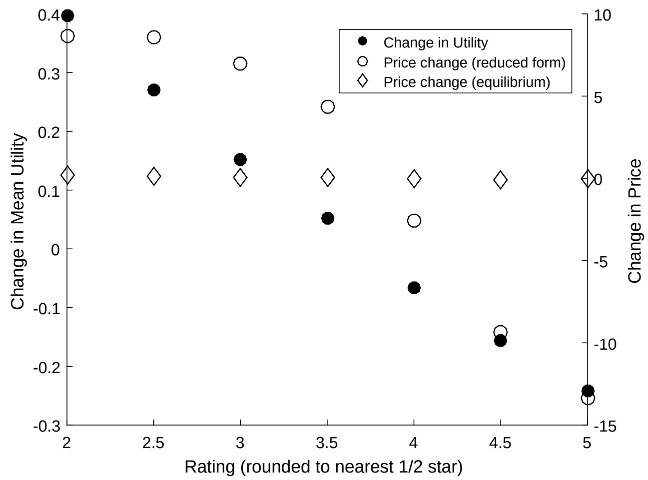

Figure 11: Inputs to the counterfactual estimation.

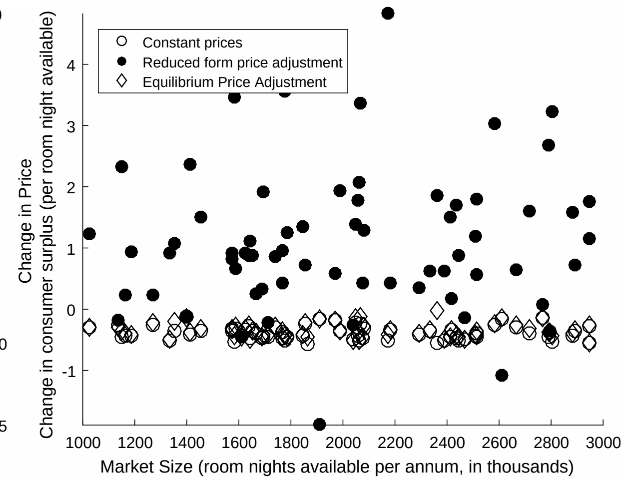

Figure 12: Changes in consumer surplus by market.

that low rated hotels are expected to increase prices in the counterfactual by up to $10 a room night, while high rated hotels drop prices by almost $15. Last, the unshaded diamonds show the difference between equilibrium prices. There is essentially no difference: the logit model predicts very little response to the counterfactual change.

Figure 12 shows the predicted changes in consumer surplus by market. In both the baseline (no price change) and equilibrium models, we see negative effects of removing information, of various sizes between $0 and -$1 per potential customer (recall that in a logit model, many customers choose the outside good). By contrast, under reduced form pricing, the counterfactual is actually generally better for consumers (and substantially so in many markets).

We dig into the differences in these results in Table 7. The top line gives the headline results: consumer surplus falls by between $123M without price adjustment, by $107M with equilibrium price adjustment, but increases by $546M in the case of reduced form price adjustment.

In all scenarios, the average percentage change in market share is positive under the counterfactual. But the effects vary with the hotel rating. Low-rated hotels gain market share, while high-rated hotels lose. In the reduced form case the strong price adjustments graphed in figure 11 mitigate the effect of the lost information, so that shares rise less at low-quality hotels (since they raise prices significantly) and fall less at low-quality hotels.

Interestingly, the effects on revenues are similar across all scenarios and for all hotel qualities (with the exception of the best hotels, who are predicted to do better under reduced form pricing). The average percentage change in hotel revenues ranges from 2.5% to 3%

Overall, the results are somewhat ambiguous. As one would expect from the theory, without any price response consumers lose when from moving to a world in which they are less informed. But the results of the reduced form pricing analysis suggest that the price responses may in fact reverse this conclusion.

6 Conclusion

It is common wisdom that consumers are devoting increasing attention and time to making informed choices, accessing the many information sources at their disposal. This has implications for market structure. In this paper, we have documented the increasing correlation between online reviews and hotel bookings, and shown some evidence that popular platforms such as TripAdvisor have causal effects on purchasing behavior. We have also demonstrated that on the supply-side, hotels have responded to this phenomenon, with high-rated hotels increasing their prices, while low-rated hotels drop theirs. Online ratings are having real effects on these markets.

This has had welfare implications. We estimate that travelers to these markets enjoy increased surplus of around \$124 million as a result of such reviews, at least when holding prices fixed. But allowing for price adjustment has ambiguous effects, depending on how we model such adjustment. It appears plausible that consumer surplus may actually be higher in the counterfactual scenario, as good hotels become unable to exploit their perceived quality through pricing.

There are many interesting directions for future research. We have not investigated the text content of online reviews, which may play an important role in matching consumers to hotels and is thus another source of surplus. Modeling information acquisition may also be important, as online reviews have almost certainly had an effect on the amount of time that consumers spend acquiring information.

References

M. Anderson and J. Magruder. Learning from the crowd: Regression discontinuity estimates of the effects of an online review database. *The Economic Journal*, 122(563):957–989, 2012.

J. Bai. Melons as lemons: Asymmetric information, consumer learning and seller reputation. 2015.

J. R. Brown, A. Goolsbee, et al. Does the internet make markets more competitive? evidence from the life insurance industry. *Journal of Political Economy*, 110(3):481–507, 2002.

E. Brynjolfsson and J. Oh. The attention economy: measuring the value of free digital services on the internet. 2012.

E. Brynjolfsson, Y. Hu, and M. D. Smith. Consumer surplus in the digital economy: Estimating the value of increased product variety at online booksellers. *Management Science*, 49(11):1580–1596, 2003.

L. Cabral and A. Hortacsu. The dynamics of seller reputation: Evidence from ebay. *The Journal of Industrial Economics*, 58(1):54–78, 2010.

J. A. Chevalier and D. Mayzlin. The effect of word of mouth on sales: Online book reviews. *Journal of marketing research*, 43(3):345–354, 2006.

D. Elfenbein, R. Fisman, and B. McManus. Market structure, reputation, and the value of quality certification. Technical report, National Bureau of Economic Research, 2014.

A. Fradkin, E. Grewal, D. Holtz, and M. Pearson. Reporting bias and reciprocity in online reviews: Evidence from field experiments on airbnb. 2014. Cited with permission.

A. Goolsbee, P. J. Klenow, et al. Valuing consumer products by the time spent using them: An application to the internet. *American Economic Review*, 96(2):108–113, 2006.

N. Hu, J. Zhang, and P. A. Pavlou. Overcoming the j-shaped distribution of product reviews. *Communications of the ACM*, 52(10):144–147, 2009.

G. Z. Jin and P. Leslie. The effect of information on product quality: Evidence from restaurant hygiene grade cards. *The Quarterly Journal of Economics*, 118(2):409–451, 2003.

S. Koulayev. Search for differentiated products: identification and estimation. *The RAND Journal of Economics*, 45(3):553–575, 2014.

M. Luca. Reviews, reputation, and revenue: The case of Yelp.com. *Harvard Business School NOM Unit Working Paper*, (12-016), 2011.

D. Mayzlin, Y. Dover, and J. Chevalier. Promotional reviews: An empirical investigation of online review manipulation. *The American Economic Review*, 104(8):2421–2455, 2014.

S. Pantea and B. Martens. The value of the internet for consumers. *Available at SSRN 2446962*, 2014.

Table 1: Demand Regressions

	(1) Log Share Ratio	(2) Log Share Ratio	(3) Log Share Ratio	(4) Log Share Ratio
Avg. Rating	0.065***			
	(0.00)			
Chain Management × Avg. Rating		0.043***		
		(0.01)		
Franchise × Avg. Rating		0.063***		
		(0.01)		
Independent × Avg. Rating		0.104***		
		(0.02)		
Economy × Avg. Rating			0.029***	
			(0.01)	
Luxury × Avg. Rating			0.249**	
			(0.08)	
Midscale × Avg. Rating			0.053***	
			(0.01)	
Upper Midscale × Avg. Rating			0.057***	
			(0.01)	
Upper Upscale × Avg. Rating			0.148***	
			(0.02)	
Upscale × Avg. Rating			0.103***	
			(0.02)	
2005 × Avg. Rating				0.010
				(0.01)
2006 × Avg. Rating				0.036***
				(0.01)
2007 × Avg. Rating				0.073***
				(0.01)
2008 × Avg. Rating				0.100***
				(0.01)
2009 × Avg. Rating				0.051***
				(0.01)
2010 × Avg. Rating				0.069***
				(0.01)
2011 × Avg. Rating				0.119***
				(0.01)
2012 × Avg. Rating				0.156***
				(0.01)
2013 × Avg. Rating				0.194***
				(0.01)
2014 × Avg. Rating				0.251***
				(0.01)
N	442,497	442,497	442,497	442,497
adjusted r-squared	0.9616	0.9617	0.9618	0.9624

Avg. rating represents the weighted cumulative average of all rating platform ratings. The dependent variable for all specifications is the log of monthly quantity rooms sold, adjusted by alpha times price. Regressions are OLS with hotel, and market × time interaction fixed effects, and errors are clustered by hotel and market × time. Specifications are 1. Rating coefficient only 2. Rating × class interaction 3. Rating × hotel ownership interaction 4. Rating × year interaction. Significance levels are denoted by asterisks (* p <0.1, ** p <0.05, *** p <0.01).

Table 2: Share Regressions

	(1) Log Share Ratio	(2) Log Share Ratio	(3) Log Share Ratio	(4) Log Share Ratio
Avg. Rating	0.088***			
	(0.01)			
xi	1.008***			
	(0.01)			
Chain Management × Avg. Rating		0.094***		
		(0.01)		
Franchise × Avg. Rating		0.080***		
		(0.01)		
Independent × Avg. Rating		0.113***		
		(0.01)		
Chain Management × xi		1.015***		
		(0.01)		
Franchise × xi		0.989***		
		(0.01)		
Independent × xi		1.039***		
		(0.02)		
Economy × Avg. Rating			0.034***	
			(0.01)	
Luxury × Avg. Rating			0.183***	
			(0.01)	
Midscale × Avg. Rating			0.061***	
			(0.01)	
Upper Midscale × Avg. Rating			0.076***	
			(0.01)	
Upper Upscale × Avg. Rating			0.134***	
			(0.01)	
Upscale × Avg. Rating			0.100***	
			(0.01)	
Economy × xi			0.902***	
			(0.01)	
Luxury × xi			1.085***	
			(0.04)	
Midscale × xi			0.913***	
			(0.01)	
Upper Midscale × xi			0.899***	
			(0.02)	
Upper Upscale × xi			0.948***	
			(0.03)	
Upscale × xi			0.906***	
			(0.02)	

2005 × Avg. Rating				0.001
				(0.00)
2006 × Avg. Rating				0.030***
				(0.01)
2007 × Avg. Rating				0.062***
				(0.01)
2008 × Avg. Rating				0.081***
				(0.01)
2009 × Avg. Rating				0.094***
				(0.01)
2010 × Avg. Rating				0.101***
				(0.01)
2011 × Avg. Rating				0.139***
				(0.01)
2012 × Avg. Rating				0.170***
				(0.02)
2013 × Avg. Rating				0.200***
				(0.02)
2014 × Avg. Rating				0.249***
				(0.02)
2005 × xi				0.999***
				(0.00)
2006 × xi				1.015***
				(0.01)
2007 × xi				1.039***
				(0.01)
2008 × xi				1.052***
				(0.01)
2009 × xi				0.976***
				(0.01)
2010 × xi				0.974***
				(0.01)
2011 × xi				0.982***
				(0.02)
2012 × xi				0.987***
				(0.02)
2013 × xi				0.999***
				(0.02)
2014 × xi				1.015***
				(0.02)
N	363,964	363,964	363,964	363,964
adjusted r-squared	0.9461	0.9464	0.9487	0.9472

Avg. rating represents the weighted cumulative average of all rating platform ratings. The dependent variable for all specifications is the log of the share ratio, adjusted by alpha times price. Regressions are OLS with hotel, and market × time interaction fixed effects, and errors are clustered by hotel and market × time. Specifications are 1. Rating coefficient only 2. Rating × class interaction 3. Rating × hotel ownership interaction 4. Rating × year interaction. Significance levels are denoted by asterisks (* p <0.1, ** p <0.05, *** p <0.01).

Table 3: Price Regressions

	(1)	(2)	(3)	(4)
	Price	Price	Price	Price
Avg. Rating	1.455***			
	(0.19)			
Chain Management × Avg. Rating		0.780		
		(0.53)		
Franchise × Avg. Rating		1.249***		
		(0.20)		
Independent × Avg. Rating		3.468***		
		(1.03)		
Economy × Avg. Rating			-0.045	
			(0.21)	
Luxury × Avg. Rating			11.529**	
			(4.47)	
Midscale × Avg. Rating			1.268***	
			(0.32)	
Upper Midscale × Avg. Rating			1.245***	
			(0.37)	
Upper Upscale × Avg. Rating			4.030***	
			(1.20)	
Upscale × Avg. Rating			2.737***	
			(0.55)	
2005 × Avg. Rating				-0.929***
				(0.25)
2006 × Avg. Rating				0.898***
				(0.25)
2007 × Avg. Rating				2.800***
				(0.26)
2008 × Avg. Rating				4.016***
				(0.28)
2009 × Avg. Rating				0.244
				(0.29)
2010 × Avg. Rating				0.186
				(0.31)
2011 × Avg. Rating				1.915***
				(0.35)
2012 × Avg. Rating				3.787***
				(0.42)
2013 × Avg. Rating				6.028***
				(0.53)
2014 × Avg. Rating				9.276***
				(0.67)
N	442,497	442,497	442,497	442,497
adjusted r-squared	0.9561	0.9561	0.9563	0.9567

Avg. rating represents the weighted cumulative average of all rating platform ratings. The dependent variable for all specifications is the room price. Regressions are OLS with hotel, and market × time interaction fixed effects, and errors are clustered by hotel and market × time. Specifications are 1. Rating coefficient only 2. Rating × class interaction 3. Rating × hotel ownership interaction 4. Rating × year interaction. Significance levels are denoted by asterisks (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

Table 4: Hotel Characteristics by Rating Change

	Rating Shift	No Change
	mean/sd	mean/sd
number of rooms	135.91	131.42
	(141.22)	(131.85)
rooms filled (monthly)	3065.38	2996.85
	(3548.89)	(3433.99)
price	108.55	118.36
	(63.99)	(98.80)
avg. rating	3.88	3.70
	(0.51)	(0.96)
number of reviews at merge	284.05	241.62
	(379.34)	(468.37)
hotel age	26.48	29.45
	(18.79)	(20.84)
N	2,706	1,330

Groups are divided into hotels that experienced a change in their rating on at least one of Expedia or Hotels.com, and hotels that were unchanged on both platforms.

Table 5: Ratings Change Experiment

	(1)	(2)	(3)	(4)
	Log Q	Log Q	Price	Price
Rating change	0.267**	0.270**	5.504	5.507
	(0.13)	(0.13)	(6.09)	(6.09)
Counterfactual rating		0.067***		0.070
		(0.02)		(0.79)
N	50,856	50,856	50,856	50,856
adjusted r-squared	0.9770	0.9771	0.9733	0.9733

Ratings changes represent the difference between the new and old average in rounded rating displayed on the three sites after the Hotels.com/Expedia merger. Counterfactual rating is the average rounded rating had the merger never occurred. Specifications are 1. ratings on log of quantity rooms sold (adjusted by alpha) 2. ratings on log of quantity rooms sold (adjusted by alpha), controlled with counterfactual rating 3. ratings on price level 4. ratings on price level with counterfactual control. Regression includes observations within a six-month window of the experiment month (June 2013). Only hotels with complete data for this period are included. Significance levels are denoted by asterisks (* p <0.1, ** p <0.05, *** p <0.01).

Table 6: Experiment Placebo

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Log Q	Log Q	Log Q	Log Q	Log Q	Log Q	Log Q	Log Q	Log Q	Log Q	Log Q	Log Q
Avg. Rating Shock	-0.078	-0.011	0.171	0.130	0.153	0.270**	0.158	0.028	0.159	0.206	-0.006	0.105
	(0.10)	(0.09)	(0.12)	(0.14)	(0.15)	(0.13)	(0.11)	(0.09)	(0.13)	(0.14)	(0.15)	(0.13)
Counterfactual Avg.	0.045**	0.057***	0.070***	0.109***	0.086***	0.067***	0.056***	0.034	0.042	0.026	0.008	0.010
	(0.02)	(0.02)	(0.02)	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)	(0.03)	(0.03)	(0.03)	(0.02)
N	50,562	50,645	50,718	50,833	50,833	50,856	51,117	51,201	51,247	51,344	51,401	51,454
adjusted r-squared	0.9775	0.9768	0.9779	0.9777	0.9779	0.9771	0.9774	0.9784	0.9790	0.9794	0.9798	0.9796

Regressions are of the value of the rating shock on log of quantity rooms sold. The counterfactual (non merger) rating is also included to control. Specifications are the months of 2013 in order, with each rating shock calculated to imply the merger had occurred in the given month. Regressions include observations within a six-month window of the hypothetical experiment month. Only hotels with complete data for this period are included. Significance levels are denoted by asterisks (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

Table 7: Counterfactual Results

	Constant Prices	Reduced Form Pricing	Equilibrium Pricing
Total change in consumer surplus ($M)	-123.95	546.25	-107.21
Average percent change in market share	2.62	0.41	2.54
Average percent change in market share (rating <=2)	43.48	24.06	43.07
Average percent change in market share (rating 3)	17.13	3.69	16.91
Average percent change in market share (rating 4)	-3.54	-2.22	-3.57
Average percent change in market share (rating 5)	-15.19	1.53	-15.12
Average percent change in revenue	2.62	3.16	2.57
Average percent change in revenue (rating <=2)	43.48	50.75	43.48
Average percent change in revenue (rating 3)	17.13	15.23	17.06
Average percent change in revenue (rating 4)	-3.54	-3.29	-3.59
Average percent change in revenue (rating 5)	-15.19	-5.47	-15.18

The constant pricing specification compares outcomes assuming prices are held fixed at observed levels. The reduced form pricing specification uses predicted prices in both the current and counterfactual scenarios, where those predictions come from the earlier price regressions. The equilibrium pricing scenario uses prices computed from solving for a complete information Nash equilibrium in prices, given the current and counterfactual gross utilities perceived by consumers. Consumer surplus is measured in million dollars, and summed across all markets. For market share and revenues, we compute the percentage change for each hotel, and then average either across the whole sample, or within a particular ratings group. The ratings groups are based on average cumulative reviews, rounded to the nearest integer. The group <=2 includes hotels with no ratings.

A Appendix

A.1 Estimation of the price coefficient

In the main text, we note that the price coefficient in the demand system, α, is estimated flexibly using supply side moments and seasonal demand fluctuations. Here we detail exactly what this means. We assume that the residual demand curve faced by each hotel j takes the following form:

$$q_{j,t} = A_j X_{j,t} - b_j p_{j,t} + \varepsilon_{j,t}$$

for some demand shifters $X_{j,t}$. We assume the demand shifters are exogenous (i.e. $E[x_{j,t}\varepsilon_{j,t}] = 0$ for each column $x_{j,t}$ of $X_{j,t}$). We assume that the hotel-specific coefficients A_j and b_j are continuous functions of the average prices $\bar{p}_j = \frac{1}{T}\sum_{t=1}^{T} p_{j,t}$ and average quantity $\bar{q}_j \frac{1}{T}\sum_{t=1}^{T} p_{j,t}$ sold by the hotel. This allows for a very flexible specification of each residual demand curve, with the main constraint being that we require hotels that sell similar numbers of rooms at similar prices to face similar residual demands. Notice that marginal revenues are given by:

$$MR_{j,t} = A_j X_{j,t} - 2b_j p_{j,t} = q_{j,t} - b_j p_{j,t} - \varepsilon_{j,t}$$

We assume also that marginal costs take the form:

$$c_{j,t} = \bar{c}_j + \omega_{j,t}$$

where the mean-zero cost shocks $\omega_{j,t}$ are assumed iid. Now the first-order condition of the profit maximization problem requires that prices are set so that marginal revenues are equal to marginal costs. So we have that $q_{j,t} - b_j p_{j,t} - \varepsilon_{j,t} - \bar{c}_j = \omega_{j,t}$, where we have rearranged terms a little. Now suppose we have an instrument $z_{j,t}$ that is correlated with quantity $q_{j,t}$ but orthogonal to the supply shock $\omega_{j,t}$. We can use this to form the following moment condition:

$$E[(q_{j,t} - b_j p_{j,t} - \bar{c}_j)\, z_{j,t}] = 0$$

Under the assumptions that marginal costs are not seasonal, two such instruments are the expected quantity sold by hotel j in the high-season, $E[q_{j,t}|j, \text{High-season}]$, and the expected quantity sold by hotel j in the low-season $E[q_{j,t}|j, \text{Low-season}]$. We construct the empirical analogs to these objects as $\hat{q}_j^H = \frac{1}{T_H}\sum_{t\in H} q_{j,t}$ and $\hat{q}_j^L = \frac{1}{T_L}\sum_{t\in L} q_{j,t}$. Given these two instruments and a little more algebra, we can solve for b_j hotel-by-hotel as $\hat{b}_j = \frac{\hat{q}_j^H - \hat{q}_j^L}{\hat{p}_j^H - \hat{p}_j^L}$, where $\hat{p}_j^H$ and $\hat{p}_j^L$ are the average high-season and low-season prices for hotel j. But such estimates are going to be noisy, and so instead we pool, estimating $b(\bar{p}, \bar{q})$ pointwise on a grid of $(\bar{p}, \bar{q})$ values by weighted instrumental variables regression of price on quantity, where we instrument for quantity with the interaction of hotel identity and a dummy for high season; and we weight observations using a Gaussian product kernel in $(\bar{p}, \bar{q})$-space. This delivers slope coefficients and elasticity estimates for each hotel. As shown in the paper, these are reasonably sensible, with hotels that face more competition (in terms of similarly sized hotels, or hotels that price similarly) having more elastic residual demand.

But since in the paper our demand system is a logit, with a single parameter α governing the way price affects demand, we must project down these elasticities to a lower dimensional space. Specifically, in a logit, the demand elasticity takes the form $\alpha p_{j,t}(1 - s_{j,t})$, and so we recover an estimate of α by matching moments as described in the main text.