Competitive Responses in a Devastated Industry: Evidence from Hotels during COVID-19

Michael D. Noel*

May 27, 2020

Abstract

Little is known about how firms devastated by the COVID-19 pandemic balanced difficult business options in the face of collapsing demand. I examine one particularly hard hit industry, hotels, and investigate competitive responses along the extensive margin (the decision to close) and the intensive margin (price changes). I find significant hotel closures and price decreases from pre-pandemic expectations to mid-pandemic expectations, even holding the timing of consumption constant. Effects are highly heterogeneous, and largely depend on how well non-price product characteristics match traveler preferences in a pandemic. Prices even increased in certain cases. Local infection rates have surprisingly little effect.

JEL Classification Codes: I10, L11, L16, L83, M21

Keywords: Novel coronavirus, virus, infections, COVID-19, pandemic, social distancing, accommodations, hotels, lodging, exit regressions, conditional price effects

1 Introduction

The collapse of the U.S. economy triggered by the COVID-19 coronavirus in March 2020 was less of an economic meltdown and more of an economic flash fire. In the four week period from mid-March to mid-April, 22.1 million Americans applied for unemployment benefits and the U6 underemployment rate skyrocketed to 22.8%, approaching Great Depression era numbers (DOL (2020), BLS (2020)). Stay-at-home orders were issued, non-essential businesses were shuttered, and even essential businesses remaining open saw a collapse in demand of a scale and speed not previously imagined. Economists have only begun analyzing the macroeconomic consequences of COVID-19, and even less is known about the microeconomic decisionmaking taking place in the

*Texas Tech University, michael.noel@ttu.edu.

"trenches" of the pandemic response. In particular, how individual firms balance multiple margins of response along extensive and intensive margins to position themselves for the best chance of survival in a post-pandemic world.

Herein, I take a closer look at microeconomic responses by firms in the context of one hard hit industry, the hotel and lodging industry. The industry was devastated by the pandemic like many others but serves as a particularly interesting application for several reasons. First, it was an essential industry not required to close, and generally not required to, or not needing to, artificially restrict capacity in any way. Travel demand collapsed, hotel occupancy rates fell to an all-time low of 21.6% in early April 2020 (STR (2020)), and hotel operators were forced to weigh a largely unrestricted set of difficult options for dealing with the new reality. Second, the industry operates on an advance purchase model, where one can book the product for future consumption at a current price that fluctuates over time. This distinction between the time of purchase and the time of consumption and the availability of data on both will have significant implications for identification.

Using operations data and pricing information from over five thousand independent hotel operators affiliated with a major chain, I explore how the pandemic forced managers to rethink decisions at the extensive margin (the decision whether to remain open or close) and at the intensive margin (how to reset prices if remaining open) in response to an unprecedented collapse in demand. I ask what decisions hotel operators were making and how these decisions were shaped by the product and service characteristics of the hotels themselves. An important issue will be how product and service characteristics interact with the changing needs of travellers, either favorably or unfavorably, in the midst of a pandemic.

The questions are as simple as they are fundamental. The first set of questions addresses issues on the extensive margin. Would the pandemic cause mass closures in the hotel industry, or would most hotels remain largely open in spite of paltry occupancy rates? Which types of hotels would be most likely to close and why? Which product characteristics are relatively more valued? Would a hotel's native amenability to social distancing practices matter? Would hotel operators and their customers be wary of infection hotspots and, if so, would they be more likely to close in infection hotspots than in other areas?

The second set of questions addresses the intensive margin for those hotels remaining open. Would hotels necessarily lower their prices to attract consumers and, if so, by how much? Would any hotels actually raise prices and, if so, in which cases and why? Would price effects be heterogeneous and if so, what determines the heterogeneity? Would hotel and area characteristics matter? Would a hotel's native amenability to social distancing practices matter? Would price changes be more pronounced in infection hotspots or be similar in all areas? And how long did hotel operators and consumers expect all these effects to last?

By examining the individual decisions of firms, how they depend on the nature of the firm, the circumstances of the firm, the environment surrounding of the firm, I can shed some light on responses to the pandemic at the micro level, and reasons for them. There are many expected results and some unexpected ones. Businesses and their physical infrastructures were not designed with this type of problem in mind, and it leads to an interesting set of strongly heterogeneous responses. Responses can be affected by both hotel characteristics and area characteristics, including infection and death rates from the coronavirus. Heterogeneous responses at the micro level are not well understood but important as they ultimately aggregate to produce the macroeconomic outcomes that are the subject of so many other studies.

While the expected effect of the pandemic on overall closures is unambiguously negative, and is a question of degree and heterogeneity, the overall effect on prices is more ambiguous. Price decreases are the likely outcome, but price changes depend on firm-level elasticities, which in turn depend on the elasticity of demand of residual consumers still active in the marketplace. With leisure travel shuttered and business meetings moving online, the residual consumer still travelling and staying in hotels in the middle of a pandemic is unlikely to be the typical customer in normal times. They may have a more critical purpose or may put additional weight on previously-unappreciated non-price factors, such as a hotel's native amenability to social distancing. So while on one hand, hotels competing with record-low occupancy rates push for universally lower prices, on the other hand, a meaningful decrease in the elasticity of the residual consumer, and a shift in where residual consumers go, may decrease firm-level elasticities in some cases and lead to smaller price decreases or even higher prices in some cases. As discussed below, a similar effect occurs in the pharmaceutical industry where it is common for branded drug manufacturers to increase prices post-patent-expiry, in spite of falling firm-level demand.

Identification strategy is important in reduced-form studies, and several unique features of the current research environment is amenable to yielding especially robust estimates. First, and obviously, the onset of the pandemic was almost entirely unexpected a few months before, and can reasonably be considered a strictly exogenous variable. Second, and unique to this study, is the availability of a three-dimensional panel of data, instead of the more common two-dimensional panel or single aggregated time series. Many micro-based studies relating to the pandemic will be limited to before-and-after type comparisons, comparing outcomes after the start of the pandemic to outcomes before, as they may not have access to an unimpacted control group, such is the nature of pandemics. And while the pandemic is the obvious shock, it would be more ideal if a control group could be used, unimpacted by the pandemic, that shows what outcomes would have been after the start of the pandemic but without the pandemic. Fortunately, the advance purchase model of the hotel industry allows me to do exactly that - observe expectations of outcomes after the start of the pandemic, but as they are expected before the start of the pandemic. One can then compare them with actual outcomes after the start of the pandemic. Specifically, instead of comparing prices for hotel stays before the start of the pandemic to prices for hotel stays after, I can compare prices for the same room at the same hotel for a stay on the same date (say, in April in the midst of the pandemic), first as expected before the start of the pandemic and again as realized after the start of the pandemic. Similarly, instead of comparing hotel closures before the start of the pandemic to hotel closures after the start of the pandemic, I can compare hotel closures for the same date in April, first as expected before the start of the pandemic and again as realized after the start of the pandemic. This is only possible because of three dimensional nature of the data - the first dimension being the identity of the hotel, the second dimension being the date on which the given hotel stay is to occur, and the third dimension being the date on which a current price for a given future stay is active. This new third dimension substantially improves identification by removing unobserved but known-to-the-firm time-varying factors correlated with the onset of pandemic. It allows the researcher to isolate the effects of the pandemic on closures and prices in an especially controlled way that may not be possible in many other industries.

2 Background and Literature

The novel coronavirus COVID-19 caused a historic and sudden shutdown of the U.S. economy in mid-March 2020. Chinese officials first reported a contagious flu-like illness in Wuhan, China, on December 31st, 2019, the World Health Organization reported the first known case outside of China on January 15th, and the U.S. Centers for Disease Control and Prevention reported the first known case in the U.S. on January 20th (WHO (2020), Holshue et al. (2020)). The data shows that at the end of February, there were only nineteen known cases in the U.S., but by the end of May, there were more than 1.5 million U.S. cases and over 100,000 deaths.

The first tangible indication of stress of the U.S. economy came the week of February 24th, when the stock market began a significant period of volatility (Baker et al. (2020a)). It was triggered in part by news of community spread and by a prediction from the director of the National Center for Immunization and Respiratory Diseases that it would not be a question of 'if' the virus would spread, but how many Americans would be severely affected when it did (CNN (2020)). The virus indeed spread quickly and, by mid-March, state governments began closing down large parts of the economy. By the end of May, forty million more Americans had applied for unemployment benefits (DOL (2020)).

An ever-growing literature examines the economic impacts of the virus on a variety of macroeconomic outcomes, including stock market outcomes (Baker et al. (2020a), Alfaro et al. (2020)), labor market outcomes (Coibion et al. (2020), Baek et al. (2020)), economic uncertainty (e.g. Baker et al. (2020b), aggregate consumption and GDP (Chen et al. (2020), del Rio-Chanona (2020)), and others. Many studies examine possible recovery scenarios with or without government intervention (Eichenbaum et al. (2020), Stock (2020), Atkeson (2020), Barro et al. (2020), McKibbin & Fernando (2020)), Guerrieri et al. (2020)). While the early economic literature on COVID-19 naturally focuses on macroeconomic outcomes, very little is known about the microeconomics of firm-level decisionmaking in the face of a pandemic, and that is the focus here.

I use the hotel industry as my application. As such, this study shares a common thread with other studies in the Industrial Organization and Management literatures that also examine the hotel industry, albeit in more normal situations. These include entry studies (Mazzeo 2002, Enz et al. (2008)), and pricing studies (Abrate et al. (2012), Hung et al. (2010), Tappata & Cossa (2014), Guo et al (2013), Koulayev (2014)) in particular.

The current study has an interesting overlap with the branded pharmaceutical pricing literature. It is well known that when low-cost generic competitors begin producing a drug after the branded drug manufacturer loses patent protection, the branded firm sometimes responds by increasing, rather than decreasing, its branded price (Frank & Salkever (2004), Berndt & Conti (2018)). The reason is that residual consumers who insist on buying the branded drug even when much cheaper generics are available tend to be more inelastic than the typical customer. If there are enough such consumers, the branded firm could profitably raise prices on them, even with lower firm-level demand. In the current context, firm demand is lower not because of new entry but because of a pandemic, yet the same logic applies. If residual consumers travelling in a pandemic in spite of widespread stay-at-home orders are more inelastic than the typical customer, some hotels may be able to raise prices in spite of lower demand and offset some of their losses.

3 Data

The study utilizes data from hotel operations, including open/closed status, prices, locations and other hotel characteristics, and characteristics of the area in which they are located, including population demographics and COVID-19 infection and death rates.

The hotel operations data derives from a major hotel chain that has over five thousand hotels in its network in the United States and is present in every state. The vast majority of hotels are owned or operated not by the chain but by individual hotel operators who pay the chain a franchise fee, as a percentage of revenues, for use of its brand name and inclusion in its reservations system.[1] The hotel data includes two key pieces of information - a hotel's expected open/closed status for a given set of future dates and, if open, the lowest available price for a standard room through direct booking on those dates. A closed status specifically means that a hotel is not accepting reservations for any stays over a range of dates, and does not necessarily mean permanent closure. Status and room price information were obtained for a total of 5,253 hotels for stays in the months of January,

[1] The chain manages 16% of hotels and owns only 0.5%.

April, and July 2020.

Importantly, the data was collected twice, creating a three dimensional panel. It was first collected in the last week of December 2019 before the start of the pandemic (and before there was even an inkling of one) and then again in the last week of March 2020 after the start of the pandemic (and with stay-at-home orders for non-essential travel widely in place). This first snapshot shows which hotels were expecting to be open in April (and July) after the start of the pandemic, and what prices they were expecting to get, as of late December. The second snapshot shows which hotels were still expecting to be open in April (and July), and what prices they were now expected to get, as of late March. The change in expectations from late December to late March, for the same stay on the same date in April or July, is attributable to the pandemic and its resulting collapse in demand.

I collect data on hotel characteristics, most notably on each hotel's category rating. The category rating is a summary measure of the overall quality and desirability of a hotel, including its destination quality and the level of services and personal attention it provides.[2] There are eight categories in all, with higher category ratings correspond to higher-quality, more destination-type hotels. Category 1 and 2 hotels are simple, limited-service hotels, usually located along highways outside of major cities and providing basic accommodations to largely drive-through traffic. Category 3 and 4 hotels are higher quality hotels with larger lobbies and a wider array of services (e.g. restaurants, pools, business centers), often located in major cities and popular destination areas. Category 5 and 6 hotels are large full-service hotels with upscale lobbies and personalized services, generally located in the most popular business and tourist destinations. Category 7 and 8 are the elite luxury resort hotels in the chain, few in number, and located in the heart of major metropolitan business districts and the most prestigious tourist vacation destinations.

In addition to a hotel's category rating, I collect address location information for each hotel and map it to the county in which it is located. I match each hotel with its surrounding area characteristics. I use demographic information on population and gender and race composition at the county level as reported the Census Bureau using 2015 estimates, and information on COVID-19 infections and deaths at the county level as assembled by the New York Times from state health

[2] The category rating is used to establish free night award redemption rates under its loyalty rewards program.

authorities.

Because two of the three dimensions in the data are time-based, it is important to align contemporaneous infection and death count data with the correct time dimension. In the main analysis, I align infection and death counts to be contemporaneous with the dates that expected prices for future stays were current (i.e. the third dimension of the data) rather than the dates the future stay would actually occur (the second dimension of the data), since price expectations are based on information known at the time and not ex-post realizations of what ultimately occurs. In later analyses, I explore alternate types of infections forecasts.

4 Methodology

The methodology takes advantage of the three-dimensional panel. With multiple snapshots of each hotel's April (and July) price calendar taken at different times, once before the start of the pandemic and once after, I can compare expected and "realized" closures - not for different dates - but for the same hotel on the same April date, once as expected in late December prior to the pandemic and once as "realized" in late March after the start of the pandemic. Similarly, I can compare expected prices for the same room and hotel on the same April date, once as expected in late December prior to the pandemic, and once as "realized" in late March during the pandemic. The identity of the hotel and the exact date of the stay are held constant in the comparisons, and combined with the fact that the pandemic is an exogenous shock, the empirical strategy takes the form of a natural experiment. Specifically, it calculates 1) the difference in realized April outcomes and realized January outcomes for a given hotel (with the pandemic treatment), and 2) the difference in expected April outcomes and realized January comes for that exact same hotel (without the pandemic treatment). This collapses to the difference between realized April prices (with the pandemic) and pre-pandemic expected April prices (without the pandemic), all for the same hotel and same date. Provided that pre-pandemic expectations are unbiased as expected, it is like applying a treatment to a subject and checking the response, then reversing back in time and giving a placebo to the exact same subject again and checking the response a second time.

A word on nomenclature. I will refer to expectations about April closures and prices formed in late December (as given by the late December data) as "expectations" and updated expectations about April closures and prices from late March (as given by the late March data) as "realizations". The idea is that April demand expectations formed in the last few days of March will be much closer to realized demand than those formed in late December. The terminology has no effect on the empirics, but significantly shortens explanations in the text. The term is most accurate for realized April closures since these are largely set in stone by the end of March. And while the term "realized April prices" can be taken to mean late-March expected prices for unsold April rooms, pre-booked reservations can be cancelled and rebooked without penalty or vacancy concerns, so they largely represent late-March expected prices for all April rooms.[3]

The main analysis proceeds in two stages. I first estimate a series of closure regressions, comparing each hotel's expected closure status for the month of April 2020 (as expected in December 2019 prior to the start of the pandemic) with its realized opening status for the month of April 2020 (as realized in late March 2020 after the start of the pandemic). I estimate overall closure rates and examine how these rates vary across hotel and area characteristics, including local infection rates. The basic estimating equation is:

$$\begin{aligned} E(CLOSED_{ijrt} \mid H, D, V, R, \Theta) \quad &= \quad G(\alpha^c + \beta^c H_{ijr} + \gamma^c D_{ijr} + \delta^c V_{ijr} + \rho^c R_r \\ &\quad + \phi^c (R_r * H_{ijr}) + \psi^c (R_r * D_{ijr}) + \lambda^c (R_r * V_{ijr})) \end{aligned}$$

where $CLOSED_{ijrt}$ is a dichotomous variable equal to one if hotel i located in county j will be closed at time t according to the hotel's expectation at time $r < t$, where r is either late December 2019 or late March 2020, and time t is a date usually in April. The matrix H consists of hotel characteristics (in particular its category type), the matrix D consists of area characteristics unrelated to the pandemic (population), and the matrix V consists of local infection and death rates due to the pandemic. The variable R is an indicator function equal to one for information current as of late March 2020, and zero for information current as of December 2019. I call this variable $REALIZED$ in the tables for better readability. The Θ is shorthand for all model parameters.

[3] No restrictions were placed on April cancellations due to the pandemic.

The key variables of interest are the interaction terms, which show the effect of each explanatory variable on expected closures from hotels' late December expectations to late March realizations. The data is collapsed to the monthly level in the closure regressions since all right side variables are month-of-stay-invariant except R.[4] For partial month closures, I set $CLOSED = 1$ if a hotel reported to be closed on a majority of days.[5]

Two forms of the G function are used - the main analysis uses $G(\delta) = exp(\delta)/(1 + exp(\delta))$ to produce a logit model, and the alternate analysis uses the identity $G(\delta) = \delta$ to produce a linear probability model (LPM). The LPM is sometimes disfavored, but is a useful alternate lens for viewing results in this context since pre-pandemic closures were rare and odds ratios can grow very large with small denominators. While t-statistics in an LPM are to be viewed with caution, I note up-front that statistical significance patterns across logit and LPM models turn out to be very similar, errors are all adjusted for heteroskedasticity, and 99.7% of all LPM predicted values lies within the unit line. Standard errors are clustered at the hotel level in all models.

In the second stage, I estimate a series of conditional price regressions, conditional on a hotel remaining open during the pandemic, at the daily level. I compare each hotel's expected price for a stay during the month of April (as expected in late December before the start of the pandemic) with its realized price for the same stay in the month of April (as realized in late March after the start of the pandemic). I estimate overall price changes and examine how these changes vary across hotel and area characteristics, including local infection rates. The estimating equation is given:

$$\begin{aligned} f(PRICE_{ijrt}) &= \alpha^p + \beta^p H_{ijr} + \gamma^p D_{ijr} + \delta^p V_{ijr} + \rho^p R_r \\ &\quad + \phi^p (R_r * H_{ijr}) + \psi^p (R_r * D_{ijr}) \\ &\quad + \lambda^p (R_r * V_{ijr}) + \varepsilon^p_{ijrt} \end{aligned}$$

where $PRICE_{ijrt}$ is a continuous variable representing the best available price for a standard room

[4] There will be two datapoints per hotel for the month of April, one for expected April closure ($R = 0$) and one for realized April closure ($R = 1$).

[5] The vast majority of hotels either expected to remain open or remain closed for the entire month of April. Using a percentage closure measure instead of a dichotomous one produces very similar results.

in hotel i in county j for a stay occurring at time t according to the hotel's expectations of demand at time $r > t$. Two forms of the f function is used, the identify $f(\delta) = \delta$ for a constant linear effects model and the log function $f(\delta) = \ln(\delta)$ for a constant percentage effects model. Matrices are as defined above, and the ε_{ijrt} is an i.i.d. normally distributed error term. Standard errors are clustered at the hotel level in all models.

Summary statistics are shown in Table 1. The table shows that over 90% of hotels fall in Categories 2 through 5, with the most common categories being Category 3, accounting for 36.7% of all hotels, and Category 4, accounting for 26.3% of all hotels. There are roughly equal numbers of Category 2 (13.3%) and Category 5 hotels (14.0%), with relatively fewer Category 6 hotels (6.4%), and few Category 1 (0.7%), Category 7 (1.9%), and Category 8 hotels (0.6%).

All regressions use pairwise-complete data to eliminate the potential for composition bias. In other words, hotels must be open and publishing prices at both $r = 0$ (late December) and $r = 1$ (late March) for stays at time t (in April) to be included in the conditional price regressions. I attribute any differences I find in closure status, and in the expected prices of hotels that remain open, from late December expectations ($r = 0$) to late March realizations ($r = 1$), to the pandemic and its resulting collapse in demand. The basic estimation framework is admittedly simple, absent the usual problems with non-experimental data, but that is part of its strength.

The resulting price effects are conditional price effects, conditional on hotels remaining open, and measure the prices that consumers actually paid among their still-available choices, and the prices that still-open hotels actually received. This is distinct from unconditional price effects, which measure price changes hotels would have implemented if they were all required to remain open and continue offering rooms for sale, which was not the case. I am primarily interested in conditional price effects, but later I estimate a selection model that extracts unconditional price effect estimates as well, and explore any differences between the two. Ex ante, I expect conditional and unconditional estimates to be similar, since a large majority of hotels remained open even during the pandemic.

The available data also provides some insight into early expectations about the possible length of the economic slowdown. The analysis so far focused on stays in the month of April, immediately after the start of the pandemic, but the same analysis can be repeated for stays in the month of July as well. By comparing late December expectations with late March expectations of July demand, I can ask whether hotels, at least as of late March, were still expecting a high probability of being open in July (and accepting new reservations) and, if so, how their expectations of July prices may have changed. The analysis follows the main results.

5 Results

5.1 Closures

I begin by examining the impact of the pandemic on hotel closures, first overall and then differentially by hotel and area characteristics. Table 4 presents high-level closure regressions showing the overall impact of the pandemic on hotel closures. I simply regress a hotel's closure status $CLOSED$ on the $REALIZED$ variable and a constant to estimate the overall effect. Specification (1) is based on a logit model and displays odds ratios, Specification (2) is based on a linear probability model and displays probability point changes.

Specification (1) shows that a hotel was 5.2 times more likely to report being closed in April, from its late December expectation to its late March realization.[6] Specification (2) translates this figure into probability point changes, and shows that the probability of an April closure increased by 11.3 percentage points from late December to late March. In other words, 11.3% of the hotels that had expected to be open in April reversed course and closed once the pandemic quashed travel demand. All coefficients are statistically significant (different from one in the logit specification and different from zero in the LPM) at better than the 1% level. The results of both specifications agree that the impact of the pandemic on hotel closures was severe.[7]

The odds ratio in Specification (1) is less than infinity, which may prompt the question as to why there were any hotels planning to be closed in April even as of the previous December. The constant term in the LPM regression of Specification (2) shows about 3.2% of hotels fall in this category. An investigation shows that few of these were closing or leaving the chain, and the vast majority were new hotels, either physically new buildings under construction or buildings undergoing renovations,

[6] More accurately, the odds of April closing versus April opening, according to the late March realization, was 5.2 times larger the odds of April closing versus April opening, according to the late December expectation.

[7] It is unknown how many hotel operators received federal relief dollars under the CARES Act of follow-up legislation, but it is reasonable to imagine that the closure rate would be higher without.

and joining the chain. They are in the database because they were accepting reservations for a later date, even if not currently open. It prompts the question of whether new hotels and existing hotels were likely to differ in their pandemic response in some way. On one hand, new hotels may find it easier to postpone openings, since they need not lay off as many permanent staff or cancel as many reservations, but on the other hand, new hotels tend to have higher debt obligations and may be more eager to start generating revenue.

To investigate any differential response between "existing" and "new" hotels, I separate hotels into two groups - existing hotels that were open and operating in January and new hotels that were not. Specification (3) and (4) consider existing hotels. Specification (3) is based on a logit specification and shows that an existing hotel open in January was 54.6 times more likely to be closed in April. The coefficient is statistically significant at the 1% level. The odds ratio is high, first, because the proportion of April closures (the numerator) is high and, second, because few hotels open in January were expecting to be closed in April (the denominator). An alternate look is provided by the LPM Specification (4), which gives a coefficient on $REALIZED$ of 11.3, statistically significantly different from zero, and showing a 11.3 percentage point increase in unexpected April closures. The constant coefficient is small, at 0.002, showing few expected closures, which leads to a very high odds ratio.

Specification (5) and (6) consider new hotels instead. Specification (5), based on a logit, shows that a new hotel was only twice as likely to close in April relative to its earlier expectation, with "only" being a relative term. Comparing this with Specification (3), existing hotels were impacted much more than new hotels, in odds ratio terms. But in probability point terms, Specification (6) shows a coefficient on $REALIZED$ is 11.6, statistically indistinguishable from that of existing hotels (11.3) and all hotels overall (11.3). Basically, the same percentage of hotels in each category unexpectedly closed, but from very different initial levels. The coefficient on the constant term is 0.718, much higher than the earlier 0.002, accounting for the difference.[8]

[8] One concern would be if unexpected delays in the construction schedule - unrelated to the pandemic - might cause a few additional unexpected April closures. However, this is unlikely to be the case. Because it is problematic to accept reservations and then have to contact customers and cancel them, hotels set opening dates conservatively into the future and move them forward in stages as certainty over the construction time frames comes increases. Any such bias then goes the other way - we are unlikely to see unexpected April closures due to non-pandemic-related construction delays, but rather a small number of unexpected April *openings*. Construction itself was generally unaffected by stay-at-home orders.

Table 3 presents the full model regressions that examine the relationship between closures and hotel and area characteristics, including infection and death rates. In Specifications (1) and (2), I regress closure status on $REALIZED$, a complete set of dichotomous indicator variables for the category rating of a hotel ($CATEGORY\ 2 - CATEGORY\ 8$), with $CATEGORY\ 1$ being the omitted variable, plus all interactions between the two. Specification (1) of Table 3 reports results from a logit model and Specification (2) reports results from an LPM. To preserve space, I report only the category interaction terms plus the $REALIZED$ coefficient (relevant to the omitted Category 1) in the table.[9]

Specification (1) shows that the pandemic had significant heterogeneous effects, based on the size and destination-quality of a hotel, as measured by its category rating. Impacts were severe across the spectrum and exponentially more severe on higher-category, destination-type hotels. Category 1 hotels, basic limited service hotels along highways, had the fewest unexpected closures. The coefficient on $REALIZED$ shows that a Category 1 hotel was "only" 20% more likely close in April, relative to late December expectations. The point estimate is large but not statistically significant given the small number of Category 1 hotels in the chain. The interaction coefficients show that the impact of the pandemic grew exponentially larger with higher categories. Category 2 hotels, simple hotels with limited services in mostly drive-through areas, were 77% more likely to close in April, relative to late December expectations. Category 3 and 4 hotels, more likely to be full-service hotels, with restaurants, fitness centers, and pools, and located in relatively higher demand areas, were 3.7 times (274%) and 3.6 times (260%) as likely to close. Category 5 and 6 hotels, large full-service hotels located in popular tourist and business destinations, were 5.3 times (427%) and 12.5 times (1154%) more likely to close. Category 7 and 8 hotels, the elite luxury hotels in the chain located in the most prestigious locations, were 94.6 times (8461%) more likely and 27.01 times (1706%) more likely to close in April, relative to December expectations. Coefficients are very large and statistically significantly different from one in each case.

Specification (2) confirms the severity and heterogeneity of effects in percentage point terms. The coefficient on $REALIZED$ shows that 2.4% of all Category 1 hotels unexpectedly closed in

[9] Complete regression results, for this and all subsequent regressions, including any constant terms and main effects omitted from the tables are available with the supplementary materials provided to the journal or upon request from the author.

April, relative to late December expectations. An additional 1.1% and 3.3% of Category 2 and 3 hotels unexpectedly closed in April as well, up and above that of Category 1, for a total increase of 3.5% and 5.7%, relative to late December expectations. The individual coefficients are not statistically significantly different from zero, but the respective sums are. The remaining interaction coefficients show an additional 6.8% of Category 4 hotels, 14.6% of Category 5 hotels, 35.9% of Category 6 hotels, 49.5% of Category 7 hotels, and 67.6% of Category 8 hotels, unexpectedly closed in April, up and above that for Category 1, and all relative to late December expectations.

Hotel characteristics are likely correlated with area characteristics, so I expand the model to include area characteristics as well. Larger, higher-quality hotels tend to be located where business travellers and tourists congregate, and these are often large metropolitan areas where the risk of virus transmission may also be higher. It could be that the characteristics of the surrounding area, rather than the internal characteristics of the hotel itself, are primarily responsible for hotel closures.

To test this, Specifications (3) (logit) and (4) (LPM) add county level population data to the model. The idea is that hotels located in larger, denser areas may be more likely to close, even if they are basic low-category hotels. Specification (3) shows that this is not the case. Conditional on the type of hotel, there is no additional effect of the size of the surrounding area on closures. The odds ratio is very close to, and insignificantly different from, one. Specification (4) confirms this in probability point terms, with a marginal population effect very close to, and insignificantly different from, zero. The category interactions, on the other hand, continue to be large and statistically significant, showing hotel characteristics and not area characteristics are the primary driver of closures. Even in populated centers where the risk of infection may be high, lower category hotels were significantly more likely to stay open.[10]

Specifications (5) and (6) present the full model that accounts for the effect of virus-related infections and deaths on the probability of realized April closures. For both logit and LPM models, I regress April closure status on county-level COVID-19 cases, COVID-19 deaths, county-level

[10] In unreported results, I estimated alternate models including local demographic information on gender and race. Not surprisingly, I found no significant results with respect to demographic variables, given that hotels tend to patronized by visitors travelling to the area, rather than residents living in it.

population, a complete set of hotel category indicator variables, plus interactions between all of the above and the $REALIZED$ variable. Cases are measured in cases per thousand and deaths are measured in deaths per million to avoid very small coefficients in the table. If demand decreased disproportionately more in those areas where the rate of transmission was also high, we should expect to see more closures in those areas. If instead demand decreased more uniformly across the county, the result of widespread stay-at-home orders and federal stay-at-home guidelines, we should not see as much of a difference. I report only interaction coefficients and the $REALIZED$ coefficient in the table to preserve space.

Specification (6), based on the LPM, shows a positive coefficient on infection rates, but it is surprisingly small. The coefficient implies that for every five hundred new cases in a county of five hundred thousand people (a 0.1% infection rate), the probability of closure would increase by just 1.3%. In early April, a 0.1% infection rate was well above average, so little of the variation in closure rates can be explained by infection rates. The coefficient is statistically significant but only at the 10% level. The corresponding coefficient in the logit model of Specification (5) is positive but not statistically significant. I find no effect of local death rates on closures in either case.[11] Meanwhile, the hotel category interaction coefficients continue to be high and statistically significant, increasingly so for higher category hotels. The results are consistent with the idea that widespread stay-at-home orders affected travel and business operations broadly and not specific to areas suffering from the greatest infection rates.

The impact of the pandemic on closures was severe, and the severity grew rapidly for higher-quality destination-type hotels. Odds ratios approached triple digits and closure rates exceeded 50% in percentage point terms at the higher end of the scale. While one might have expected smaller, lower-category hotels to respond the most negatively, the results are consistent with the nature of hotels and how well they "match" the needs of the travelling public in the midst of a pandemic. High category hotels tend to be large and busy hotels with busy lobbies and well-used elevators where social distancing would be hard to achieve. They tend to offer a multitude of personal services, from restaurants and bars to serviced pool areas to valet parking, which become

[11] Although weak in both cases, the relatively larger impact of infection rates over death rates has an interesting parallel with the crime literature. That literature finds that crime is generally more responsive to the probability of being caught rather than the severity of the punishment after being caught (Chalfin & McCrary (2017)).

undesirable in a pandemic, and are likely to be unavailable anyway. They are also more likely to attract consumers that arrive by air, but because of social distancing problems in airports and airplanes, air travel was even harder hit than general travel. In contrast, low category hotels tend to be simple with few guests, few public spaces, and few of the services that travellers in a pandemic may neither want nor need. They cater to vehicular traffic, have convenient highway access, and it is generally possible to come and go with little human contact. So while all types of hotels were negatively impacted, lower category hotels are likely to fare relatively better if they have more of the non-price characteristics that residual travellers value in a pandemic. Essentially, the pandemic inverts the usual rank ordering of what is considered "high-quality" and "low-quality" hotels for many people.

5.2 Price Effects

The impact of the pandemic is felt not only in the businesses that close their doors but also in those that stay open. It causes occupancy rates to collapse, and materially affects the prices hotels can still charge. In the intuitive case that prices fall, lower prices exacerbate profitability losses. But if more inelastic residual demand allows some hotels to increase prices, even a little, it can offset profitability losses some.

I examine price impacts in a series of *conditional* price regressions, i.e. conditional on that subpopulation of hotels remaining open, and report the results in Table 4. Specification (1) uses price levels as the dependent variable and yields constant dollar estimates, and Specification (2) uses the log of price and yields constant percentage estimates. As before, these are always pairwise comparisons - I compare prices for the same room in the same hotel for the same night in April, once based on late December expectations and once based on late March realizations.

Both specifications agree that the conditional price effects are large. Specification (1) shows that hotel operators *lowered* prices for April stays by an average of \$39.24, from late December expectations before the start of the pandemic to late March realizations after the start of the pandemic. Specification (2) translates this into constant percentages and shows that hotels lowered prices by an average of 23.0%.[12] The large price decreases and large occupancy rate declines

[12] The average percentage effect is given by $exp(-0.261) - 1 = 0.23$.

(of 68.5%) combine to yield an incredible 78.5% overall decrease in hotel revenues due to the pandemic.[13]

A simple ordered logit regression (not shown in the table) shows that price decreases far outnumbered price increases, but it also shows some heterogeneity. The probability of prices falling was 0.83 (s.e. 0.004), the probability of prices remaining stable was 0.03 (s.e. 0.001) and the probability of prices rising was 0.13 (s.e. 0.003). Thus, a non-negligible proportion of prices actually increased from late December expectations to late March realizations, and I explore this in more detail later.

Specifications (3) and (4) limit the data to only existing hotels, and show that April prices fell an average of \$39.09, or 22.9%, relative to late December expectations. Specification (5) and (6) limit the data to only new hotels and show that April prices fell an average of \$61.53, or 30.4%, relative to late December expectations. All coefficients are statistically significant at the 1% level and statistically significantly different from each other in cross-equation tests. It shows a potentially meaningful difference in pricing strategy between existing and new hotels. The reason is unclear, but one plausible explanation relates to cash flow. If higher capital expenditures following a new build or renovation applies greater pressure on a new hotel to generate positive revenue flow, it can incentivize such hotels to price more aggressively. The effect persists even controlling for other factors.

Table 5 presents the full model regressions that examine the relationship between conditional price effects and hotel and area characteristics, including infection and death rates. It shows substantial heterogeneity in the distribution of price effects. Specifications (1) and (2) add hotel type to the model and regress price (or log price) on $REALIZED$, a complete set of dichotomous indicator variables for the category rating of a hotel ($CATEGORY\ 2 - CATEGORY\ 8$), with $CATEGORY\ 1$ being the omitted variable, and all interactions between the two. Specification (1) uses a linear-linear functional form and the Specification (2) uses a log-linear functional form.

The results show the impact of the pandemic on prices was severe, especially so on higher-category destination-type hotels. The coefficients on $REALIZED$ show that a Category 1 hotel

[13] April occupancy rates fell 68.5% on average year over year, and April prices fell 23.0% on average for 88.7% of hotels, and 100% for the remaining 11.3% of hotels.

discounted its April prices, from late December expectations to late March realizations, an average of \$8.29, or 8.7%. Summing relevant coefficients, Category 2 hotels discounted April prices an average of \$18.93 or 15.9%, Category 3 hotels discounted April prices an average of \$30.29, or 20.9%, and Category 4 hotels discounted April prices an average of \$45.15, or 25.5%.[14] At the higher end of the spectrum, Category 5 hotels discounted prices an average of \$64.72 (29.8%), Category 6 hotels discounted prices an average of \$83.26 (31.8%), and Category 7 hotels discounted prices an average of \$94.97 (28.3%). Category 8 hotels discounted prices an average of \$119.18, or 23.4%. All estimates are statistically significantly different from zero. The results show that the conditional price effects were large, and largest for hotels at the higher end of the spectrum, in both absolute and percentage terms.

Hotel characteristics may be correlated with area characteristics, so Specifications (3) and (4) add county population to the model. I test whether the size of the surrounding area, rather than the characteristics of the hotel itself, drives price effects. I find that it does not. The point estimates population are negative as expected, but small, and statistically significant only in Specification (3). That coefficient implies that every additional one million people in a county results in only an additional one dollar discount. Meanwhile, the hotel category interaction coefficients continue to be large and statistically significant in every case.

Specifications (5) and (6) present the full model that accounts for the effect of virus-related infections and deaths, as well as population and hotel category ratings, on pandemic-induced price changes. I regress April prices on county-level COVID-19 cases, COVID-19 deaths, county-level population, a complete set of hotel category indicator variables, plus interactions between all of the above and the $REALIZED$ variable. The results show that local area infection and death rates have little effect on pricing. Specification (5) shows a statistically significant but economically small impact of infections on price discounts, with each additional five hundred infections in a county of five hundred thousand people resulting in an only a \$1.32 discount. The corresponding coefficient in Specification (6) is not statistically significant. Death rates have no statistically significant impact on price changes in either specification, conditional of the number of infections.

[14] For example, the Category 2 absolute discount is \$8.29 + \$10.64 = \$18.93, and its percentage discount is $exp(-0.091 - 0.083) - 1 = 0.159$.

Like the closures analysis, the primary factor driving heterogeneity in the price change distribution is hotel type, as measured by its category rating. Small, limited-service hotels along highways and catering to drive-through traffic responded with the smallest average price discounts overall, while large full-service hotels that are destinations unto themselves responded with the largest average price discounts, in absolute and percentage terms. There is some evidence that infection hotspots matter, but the estimates are small and statistical significance is mixed.

The ordered logit model mentioned above shows a meaningfully large number of April price increases from late December expectations to late March realizations, in spite of severe and unexpected occupancy losses across the board. It suggests that some hotels may have found it more profitable to increase, rather than decrease, prices charged during the pandemic. It would be interesting to explore which ones and why. I discussed that a price increase in the face of declining demand would be profitable if there were a decrease in the average elasticity of demand facing the firm, which could occur if more inelastic residual consumers (still travelling in the midst of a pandemic in spite of widespread stay-at-home orders) stayed in certain hotels. I already discussed advantages that lower category hotels would have in a pandemic, namely that social distancing is more possible, and that this may attract a more price inelastic consumer than normal. We have already seen that lower category hotels were less likely to close down during the pandemic and that they tended to have smaller average discounts. If this logic holds, we should also expect to see lower category hotels increasing April prices in absolute terms the most often, from late December expectations to late March realizations, in response to a realized shift in the inelasticity of the residual consumer.[15]

A breakdown of the direction of price changes by hotel type confirms this to be true. Category 1 hotels increased prices instead of decreased them the most, 26.7% of the time, in spite of lower demand. Category 2 hotels increased prices the second most, 18.8% of the time and Category 3 hotels increased them the third most, 13.8% of the time. Percentages fall monotonically as category rating increases, up to Category 8 hotels which increased them just 1.2% of the time. The results are consistent with the idea that residual consumers in a pandemic are likely more inelastic and likely

[15] An increase in variable costs, e.g. from additional sanitizing, is an unlikely cause of price increases, given the size of pre-pandemic margins, very low occupancy rates, an oversupply of inexpensive labor, and the closure of most public areas inside hotels.

to gravitate to lower category hotels, which because of their smaller size and limited service nature, provides a better match between hotel characteristics and the preferences of the still-travelling public.

5.3 Alternate Specifications

The main results show that the pandemic had severe consequences in terms of closures and price effects, especially for higher-quality destination-type hotels. I found much heterogeneity in the effects, which are explained by hotel characteristics, and much less so by area characteristics. Infection and death counts in particular had little additional impact, consistent with the fact that the decrease in demand was widespread and not just specific to impacted areas.

Table 6 reports results from a set of alternate regressions to check the robustness of the main results. One possible concern is that hotels may base decisions on local infection and death growth rates, rather than infection and death counts. Specifications (1) and (2) revisit the main model using infection and death growth rates instead. Specification (1) is a closure regression based on the logit model and Specification (2) is the corresponding price regression based on a linear model.[16] The dependent variables differ across specifications so the point estimates are not comparable, but the patterns in the coefficients are.

The results confirm the main results presented earlier. The pandemic had a severe impact on closures and prices overall, and the heterogeneity in effects is largely driven by hotel characteristics and not area characteristics. The coefficient on the infection growth rate is statistically significant in the price regression, but economically small, and the other COVID–19 coefficients are not statistically significant. Meanwhile, the category interaction coefficients continue to be large and statistically significant, increasingly so for higher category hotels.

Another concern, and specific to the price regressions, is that prices for some April stays could

[16] There is a challenge in measuring growth rates since, as of late March, many areas had no infections and deaths, or few, leading to either zero or infinite growth rates. Also, growth rates and counts up to late March are highly correlated because infections and deaths generally started from a point of zero in the last few weeks of March. So I consider a hypothetical world in which hotel operators with rational expectations are able to predict the level of infections and deaths in their respective counties by the end of April, from current information. I calculate the infection growth rate as the monthly change in infections per thousand people and the death growth rate as the monthly change in deaths per million people. These are more easily predicted when infection and death rates, conditional on social distancing and other preventative practices, are relatively constant.

rise from late December to late March simply because of expiring advance purchase discounts. April stays, in particular early April stays, may be more expensive, biasing price effects towards zero.

This is unlikely to be a concern for several reasons. First, advance purchase discounts have limited availability and, second, they have a tendency to be replaced with similarly discounted rates once they expire. Third, they are dwarfed in magnitude by the price effects estimated here, and fourth, the bias would only go the wrong way anyway. Nonetheless, it is possible to test for any effects of advance purchase discounts on the results. Since the vast majority of discounts expire within fourteen days of the stay, early April stays are potentially impacted but late April stays are generally not. I define $ADVANCE = 1$ if a stay takes place in the last half of April and zero if it takes place in the first half. I add both it and its interaction with $REALIZED$ to the full model. If expiring discounts are materially affecting estimates, the interaction term should be negative and large, showing that early April prices are inflated and price effects understated. I find this not to be the case. Specification (3) shows that the coefficient on the $ADVANCE$ interaction is very small and not statistically significant. The point estimate corresponds to only a 37 *cent* additional discount for late April stays than early April stays, relative to late December expectations.

Another potential concern, also specific to the price regressions, is that serial correlation in prices could lead to overrejection bias. Bertrand, Duflo, and Mullainathan (2004) show in Monte Carlo simulations that standard OLS implementation in the presence of serially correlated data without standard error corrections can reject the null hypothesis of no effect (at the 5% level) almost 50% of the time when the null is true. I implement the two corrections that Bertrand et al. find work best. First, I use an arbitrary variance-covariance matrix to estimate standard errors, i.e. clustering, to account for serial correlation in the error term. This correction is not new but is already embedded into all the specifications contained in the study. Bertrand et al. show the adjustment largely cures the issue when there are many clusters - with fifty clusters, the rejection rate falls to 6% (instead of 50%) when the null is true and the 5% significance level is used. The dataset in this study has over five thousand clusters.

Second, I follow Bertrand et al. and aggregate the data up to a coarser unit of time, essentially collapsing serially correlated observations into a single observation, and removing much of the time dimension from the analysis. This approach reduces Type I error to the correct level, even with few clusters, but at the cost of increasing Type II error. I report results from the aggregated model in Specification (3). Even though the number of observations is reduced from over a quarter of a million observations to less than ten thousand, point estimates and statistical significance levels are similar to that from the full set of observations. Note that the closure regressions already natively include this correction, since the data in those regressions are already aggregated to the monthly level (with two datapoints per hotel based on $r = 0$ and $r = 1$). I do not duplicate those specifications in the table here.

5.4 Conditional and Unconditional Price Effects

A remaining question surrounds the difference between *conditional* price effects, i.e. average price changes conditional on a hotel actually remaining open, and *unconditional* price effects, i.e. average price changes in a hypothetical world in which all hotels would be required to remain open during the pandemic. I am primarily interested in conditional price effects, as these were realized, but it would be interesting to see how the two sets compare. I would expect any differences to be small, since more than 85% of hotels report April prices.

Estimating unconditional effects requires a selection equation that includes variables determining whether a hotel is likely to remain open (observed) or to close (unobserved), but that does not affect the price discount they would offer if open. The challenge in doing so is that the pandemic shocked hotel demand, and demand factors that affect hotel closures are likely to be the same as those that affect hotel prices.

One idea focuses on the length of commitment between a hotel and its customers. Most hotels cater to very short visits, (e.g. overnight, weekend, a vacation week), but three brands in the chain cater to long term stays instead. These long-stay hotels make up 25% of all hotels. They provide kitchenettes in all rooms, appliances and dishware, a small living and dining area, and a casual lobby with seating areas that is meant to replicate home living spaces. They are open to all consumers, but are actively marketed as temporary or semi-permanent living quarters for contract workers, businesspeople, and other residents. The average length of stay in these hotels is longer, approximately half of guests stay more than a week and half of those more than a month. The longer time commitment between a long-stay hotel and its existing long term guests should make these hotels less likely to close even with the onset of the pandemic, and even as its occupancy rate falls.

To test this, I define $STAY$ to be equal to one if a hotel is part of the one of these long-term-stay brands, and zero otherwise. I add both it and its interaction with $REALIZED$ to the main logit model for hotel closures and re-estimate the model in Specification (5). I find that long-term stay hotels were indeed less likely to close. The coefficient on the $STAY$ interaction variable is 0.288, less than one, and statistically significantly so. Other coefficients in the model are not meaningfully affected. While $STAY$ affects closure probabilities, I expect that it is unlikely to affect the size of price discounts in the absence of selection. Both long-stay and non-long-stay hotels competed for the same residual short-stay consumers before the start of the pandemic, both faced collapses in occupancy rates, and both continue to compete for the same short-stay consumers after the start of the pandemic. Prices are similar for both types of hotels conditional on category rating.

I use $STAY$ to estimate a two-stage Heckman selection model that includes the complete set of right hand side variables plus $STAY$ and its interaction term as selection variables in the first stage. The first-stage selection equation is based on a probit model and the second-stage price regression is based on a linear functional form.[17] Specification (6) presents the second-stage results and the resulting unconditional price effects. As expected, I find relatively little difference between the two sets of price effects and all the main conclusions carry through. The point estimates are marginally smaller than in the conditional price regressions, which is consistent with the expectation that hotels electing to close down are also those that could not profitably lower prices by the amount needed to effectively compete with others. The main driver of heterogeneity continues to be hotel characteristics, with higher category hotels offering larger discounts, in absolute and percentage terms, to compete for a residual consumer that sees them relatively less favorably during a pandemic. Coefficients on population and infection rates continue to be statistically significant, but small.

[17] Coefficients in the first-stage probit model are not easily interpreted and are not shown, but have similar significance patterns to the corresponding logit model. The $STAY$ interaction is statistically significant and positive (since these hotels are more likely to be open and their prices thus observable), and the inverse mills ratio is statistically significant and negative.

5.5 Dynamic Effects

Finally, the data includes expected price and closure information for the month of July, so it is possible to have an early look at how longer term expectations may have been affected by the pandemic. I can compare closure plans and expected prices not for April stays but for July stays, first based on late December expectations before the start of the pandemic and then again based on late March expectations after the start of the pandemic. Since it is presumably more costly to prematurely announce a closure for July only to recant the decision later, or to offer low July prices early on only to regret it later, I would expect firms to be more conservative. Changes I find would have to reflect a reasonably strong change in expectations about July prospects as of late March.

I present results in Table 7. Specifications (1) and (2) examine the change in July closures, from late December expectations to late March expectations. Specification (1) is the high-level closure regression showing overall effects and Specification (2) is the full model including hotel and area characteristics. Specification (1) yields a coefficient on the $REALIZED$ variable of 0.006, statistically significant but very small, and showing that very few hotels had halted July reservations from late December expectations to late March expectations. Recall that closure here specifically means that a hotel is not accepting new July reservations, and not necessarily that it was permanently closed. Specification (2) presents the full model and finds some heterogeneity across hotel type. At first glance it may appear that several interaction effects are large, but the coefficients on the interaction terms are necessary to offset the negative and imprecisely estimated $REALIZED$ main effect, applicable to the few Category 1 hotels in the data. Adding the relevant interaction coefficient with $REALIZED$, the estimated closure rates are less than 1% and statistically insignificant for Categories 2, 3, 7, and 8. For mid-category hotels, the estimated closure rates are also small but statistically significant, 1.3% for Category 4, 1.5% for Category 5, and 2.6% for Category 6.[18]

Specifications (3) and (4) are the corresponding price regressions for July, comparing late December expectations to late March expectations. Specification (3) shows a coefficient on $REALIZED$ of -1.074, statistically significant but very small, and implying only a one dollar price discount built into July prices as of late March. Specification (4) presents the full price effects model and

[18] Anecdotally, I learned that a few hotels moved planned renovation projects forward to coincide with a period of expected low demand, and some others were accepting reservations again by May.

shows additional heterogeneity in price responses across hotel characteristics compared with the closure regressions. There are *positive* point estimates on the lower three category interactions (with one significant), and negative point estimates on the upper five (with four significant). The highest average price increase for July stays was \$3.14, for Category 3 hotels (summing the relevant coefficients), and the highest average price discount for July stays was \$24.81, for Category 8 hotels. The price changes are still generally small and consistent with a wait-and-see approach. However, higher-category destination-type hotels were offering some larger discounts, consistent with the idea that vacation travellers often start planning months in advance.

6 Conclusion

The COVID-19 pandemic was a demand shock of a size and speed not before seen in the United States. It began to spread rapidly throughout the U.S. in early March and, within a few weeks, state governments began shutting down large portions of the economy. This study takes a microeconomic approach and examines decisionmaking by individual firms in the midst of the pandemic. I examined how firms weighed options on the extensive and intensive margins, how these decisions were heterogeneous across firms and why, and how COVID-19 infections and deaths affected them. I focused on one of the many industries hit hard by the pandemic, the hotel and lodging industry.

I utilized a new microdataset of hotel closures and prices that had several useful advantages from an econometric point of view. One was simply its size, a sample of over five thousand hotels affiliated with a major chain, ranging from basic hotels to luxury resorts, in all fifty states. Another was its shape - a three-dimensional panel that included not only information on closures and prices for different stay dates, but also changes in closure and price expectations leading up to those dates. The three dimensional nature of the data offered an interesting window into what firms were expecting market conditions to be after the start of the pandemic, but before they ever knew there would be a pandemic. By controlling for the identity of the hotel and the exact date of the stay, with only expectations changing from a strictly exogenous shock, I could estimate the impact of the pandemic on microeconomic outcomes and all in a very controlled way.

The results show that the impact on the hotel industry was severe. Average prices were 23.0%

lower for hotels that remained open, and with April occupancy rates 68.5% lower than the previous April, revenues across all hotels fell by an incredible 78.5% on average. I can gauge changes in profitability by noting that net profit margins average 8% in normal times and that fixed costs make up about 70% of total costs at normal occupancy. If variable costs are proportional to occupancy rates but fixed costs are not, total costs would fall by only 20.6%, next to a 79.3% drop in revenues. This corresponds to a 814% drop in profits overall, basically in a month. Proprietary data on occupancy rates and costs would be needed to be more specific, but revenue and profit estimates are dire under any plausible set of assumptions.

I found that a significant proportion of hotels closed down once the pandemic hit and those that remained open often reduced prices by large amounts to attract what consumers were left. The impact was very heterogeneous, with full-service destination-type hotels taking the largest brunt of the impact. Closure odds ratios approached 100-fold and price discounts exceeded 30% among the highest categories. On the other end of the spectrum, basic limited-service hotels were least likely to close and had the smallest price discounts overall. A significant number of "low-quality" hotels were actually able to increase their prices and this is consistent with a lower firm-level elasticity of demand due to more inelastic residual consumers and a shift in the non-price characteristics that residual consumers value when travelling in a pandemic. Simply put, less is more during a pandemic.

While hotel characteristics mattered a great deal, area characteristics were surprisingly less important. Population had little impact on closures and price effects, and even though there was evidence of larger effects in local infection hotspots, the effects tended to be economically small. The results are consistent with a national demand shock with little adjustment for local idiosyncrasies.

7 References

Abrate, G., Fraquelli, G., Viglia, G. (2012) “Dynamic pricing strategies: Evidence from European hotels”. *International Journal of Hospitality Management* 31(1), 160-168.

Alfaro, L., Chari, A., Greenland, A., Schott, P. (2020). “Aggregate and firm-level stock returns during pandemics, in real time”. NBER Working Paper 26950.

Atkeson, A. (2020). “What will be the economic impact of COVID-19 in the US? Rough estimates of disease scenarios.” NBER Working Paper 26867.

Baker, S., Bloom, N., Davis, S., Kost, K., Sammon, M., Viratyosin, T. (2020a). “The unprecedented stock market impact of COVID-19”. NBER Working Paper 26945.

Baker, S., Bloom, N., Davis, S., Terry, S. (2020b). “COVID induced economic uncertainty”. NBER Working Paper 26983.

Baek, C., McCrory, P., Messer, T., Mui, P. (2020). “Unemployment effects of stay-at-home orders: evidence from high frequency claims data”. IRLE Working Paper 101-20.

Barro, R., Ursua, J., Weng, J. (2020). “The coronavirus and the great influenze pandemic: lessons from the ‘Spanish Flu’ for the coronavirus’s potential effects on mortality and economic activity”. NBER Working Paper 26866.

Berndt, E., Conti R. (2018). “Specialty drug prices and utilization after loss of U.S. patent exclusivity, 2001-2007” in "Measuring and Modeling Health Care Costs", Aizcorbe A., Baker, C., Berndt, E., Cutler, D., eds. University of Chicago Press.

Bertrand, M., Duflo, E., Mullainathan, S. (2004). “How much should we trust differences-in-differences estimates?” *Quarterly Journal of Economics* 119(1), 249-275.

Bureau of Labor Statistics (BLS) (2020). “Labor Force Statistics from the Current Population Survey: U6”. Series LNU03327709.

Chalfin, A. & McCrary, J. (2017). "Criminal deterrence: A review of the literature", *Journal of Economic Literature* 55(1), 5-48.

Chen, H., Qian, W., Wen, Q. (2020). “The impact of the COVID-19 pandemic on consumption: learning from high frequency transaction data”. Working Paper.

Coibion, O., Gorodnichenko, Y., Weber, M. (2020). “Labor markets during the COVID-10 crisis: a preliminary view”. CESifo Working Paper 8238.

CNN (Mclaughlin, E. & Almasy, S.) (2020). “CDC official warns Americans it’s not a question of *if* coronavirus will spread, but when". February 26, 2020.

del Rio-Chanona, M., Mealy, P., Pichler, A., Lafond, F., Farmer, J.D. (2020) "Supply and demand shocks in the COVID-19 pandemic: an industry and occupation perspective". Working Paper.

Department of Labor (DOL) (2020). "Unemployment Insurance Weekly Claims Data". Report r539cy.

Eichenbaum, M., Rebelo, S., Trabandt, M. "The macroeconomics of pandemics". NBER Working Paper 26882.

Enz, C., Canina, A., Liu, Z. (2008). "Comparative dynamics and pricing behavior in US hotels: The role of co-location". *Scandinavian Journal of Hospitality and Tourism* 8:3, 230-250.

Frank. R., Salkever, D. (2004). "Generic entry and the pricing of pharmaceuticals". *Journal of Economics & Management Strategy* 6(1), 75-90.

Guerrieri, V., Lorenzoni, G., Straub, L., Werning, I. (2020)."Macroeconomic implications of COVID-19: Can negative supply shocks cause demand shortages?" NBER Working Paper 26918.

Guo, X., Ling, L., Dong, Y., Liang, L. (2013). "Cooperation contract in tourism supply chains: The optimal pricing strategy of hotels for cooperative third party strategic websites". *Annals of Tourism Research* 41, 20-41.

Holshue, M., DeBolt, C., Lindquist, S., Lofy, K. Wiesman, J., Bruce, H., Spitters, C., Ericson, K., Wilkerson, S., Tural, A., Diaz, G., Cohn, A., Fox, L., Patel, A., Gerber, S.,Kim, L., Tong, S., Lu, X., Lindstrom, S., Pallansch, M., Weldon, W., Biggs, H., Uyeki, T., Pillai, S. (2020). "First case of 2019 novel coronavirus in the United States". *New England Journal of Medicine* 382, 929-936.

Hung, W., Shang, J., Wang, F. (2010). "Pricing determinants in the hotel industry: Quantile regression analysis". *International Journal of Hospitality Management* 29, 378-384.

Koulayev, S. (2014) "Search for differential products: Identification and estimation". *RAND Journal of Economics* 45:3, 553-575.

Mazzeo, M. (2002). "Product choice and oligopoly market structure". *RAND Journal of Economics*

33:2, 1-22.

McKibbin, W., Fernando, R. (2020). “The global macroeconomic impacts of COVID-19: seven scenarios”. CAMA Working Paper 19/2020.

Stock, J. (2020). “Data gaps and the policy response to the novel coronavirus”. NBER Working Paper 26902.

STR (2020). "U.S. hotel results for week 4 April". Press release. April 8, 2020.

Tappata, M., Cossa, R. (2014). “Price discrimination 2.0: Opaque bookings in the hotel industry”. Working Paper.

World Health Organization (2020). "Novel Coronavirus (2019-nCOV) Situation Report - 1", January 21, 2020.

Table 1. Summary Statistics

	Num. Obs.	Mean	Std. Dev.	Min.	Max.
Price per Night	35024	171.17	91.60	23.7	2253.2
Closures	42104	0.168	0.374	0.0	1.0
Category 1	42024	0.007	0.088	0.0	1.0
Category 2	42024	0.133	0.339	0.0	1.0
Category 3	42024	0.367	0.482	0.0	1.0
Category 4	42024	0.263	0.440	0.0	1.0
Category 5	42024	0.140	0.348	0.0	1.0
Category 6	42024	0.064	0.244	0.0	1.0
Category 7	42024	0.019	0.138	0.0	1.0
Category 8	42024	0.006	0.075	0.0	1.0
County Population	42104	1136.70	1612.710	5.5	10170.3
County Infections	42104	902.61	2541.84	0.0	25271.0
County Deaths	42104	35.41	109.16	0.0	1226.1

Population in thousands. Infections and deaths in cases. Price per Night in dollars.

Table 2. High-Level Closure Regressions

Dep. Var.: Closed	(1)	(2)	(3)	(4)	(5)	(6)
Realized	5.158**	0.113**	54.585**	0.113**	1.968**	0.116**
	(22.757)	(25.395)	(13.855)	(25.122)	(4.190)	(4.266)
Constant	0.033**	0.032**	0.002**	0.002**	2.541**	0.718**
	(-43.462)	(13.131)	(-20.894)	(3.468)	(6.156)	(23.346)
Model	Logit	LPM	Logit	LPM	Logit	LPM
Jan. 2020 Status	All	All	Open	Open	Closed	Closed
Num. Obs.	10526	10526	10094	10094	432	432
Adj/Psuedo R-squared	0.071	0.040	0.164	0.057	0.018	0.017

Logit estimates expressed as odds ratios. t-statistics (LPMs) or z-scores (logit) in parentheses.
** Significant at the 5% level. * Significant at the 10% level.

Table 3. Main Closure Regressions

Dep. Var.: Closed	(1)	(2)	(3)	(4)	(5)	(6)
Realized	1.201	0.024	1.247	0.026	1.207	0.024
	(1.011)	(1.012)	(1.196)	(1.076)	(1.005)	(1.009)
Category 2 * Realized	1.765**	0.011	1.766**	0.011	1.723**	0.011
	(2.261)	(0.449)	(2.266)	(0.450)	(2.165)	(0.424)
Category 3 * Realized	3.744**	0.033	3.738**	0.033	3.619**	0.031
	(5.528)	(1.321)	(5.532)	(1.320)	(5.342)	(1.254)
Category 4 * Realized	3.600**	0.068**	3.669**	0.069**	3.492**	0.066**
	(5.607)	(2.686)	(5.706)	(2.720)	(5.377)	(2.602)
Category 5 * Realized	5.267**	0.146**	5.484**	0.148**	5.185**	0.145**
	(6.655)	(5.262)	(6.660)	(5.320)	(6.309)	(5.217)
Category 6 * Realized	12.542**	0.359**	12.940**	0.360**	12.346**	0.357**
	(8.159)	(9.990)	(8.205)	(10.028)	(8.019)	(9.944)
Category 7 * Realized	94.610**	0.495**	97.280**	0.496**	91.425**	0.491**
	(4.451)	(8.993)	(4.488)	(9.021)	(4.416)	(8.937)
Category 8 * Realized	27.061**	0.676**	27.973**	0.677**	25.727**	0.669**
	(4.816)	(7.754)	(4.843)	(7.771)	(4.670)	(7.702)
Population * Realized			1.000	0.000	1.000	0.000
			(-1.084)	(-0.735)	(-1.009)	(-0.720)
Infections/1KPop * Realiz					1.041	0.013*
					(0.364)	(1.835)
Deaths/1KPop * Realized					1.006	0.000
					(0.964)	(-1.106)
Model	Logit	LPM	Logit	LPM	Logit	LPM
Num. Obs.	10506	10506	10506	10506	10506	10506
Adj/Psuedo R-squared	0.155	0.130	0.155	0.129	0.156	0.130

Logit estimates expressed as odds ratios. t-statistics (LPMs) or z-scores (logit) in parentheses.
** Significant at the 5% level. * Significant at the 10% level.

Table 4. High-Level Price Regressions

	(1)	(2)	(3)	(4)	(5)	(6)
Realized	-39.236**	-0.261**	-39.088**	-0.261**	-61.527**	-0.363**
	(-75.377)	(-83.737)	(-75.174)	(-83.366)	(-6.849)	(-8.459)
Constant	161.786**	5.020**	161.634**	5.019**	184.609**	5.140**
	(196.547)	(1135.593)	(196.242)	(1134.137)	(14.029)	(76.064)
Dependent Variable:	*PRICE*	*ln(PRICE)*	*PRICE*	*ln(PRICE)*	*PRICE*	*ln(PRICE)*
Num. Obs.	258792	258792	257084	257084	1708	1708
Adj R-squared	0.108	0.138	0.107	0.137	0.213	0.229

t-statistics in parentheses. ** Significant at the 5% level. * Significant at the 10% level.

Table 5. Main Price Regressions

	(1)	(2)	(3)	(4)	(5)	(6)
Realized	-8.290**	-0.091**	-7.553**	-0.089**	-7.290**	-0.088**
	(-3.770)	(-3.788)	(-3.411)	(-3.715)	(-3.289)	(-3.670)
Category 2 * Realized	-10.639**	-0.083**	-10.630**	-0.083**	-10.486**	-0.082**
	(-4.583)	(-3.308)	(-4.581)	(-3.310)	(-4.516)	(-3.281)
Category 3 * Realized	-21.998**	-0.144**	-22.026**	-0.144**	-21.761**	-0.143**
	(-9.659)	(-5.888)	(-9.677)	(-5.899)	(-9.546)	(-5.851)
Category 4 * Realized	-36.863**	-0.203**	-36.456**	-0.202**	-35.997**	-0.200**
	(-15.327)	(-8.190)	(-15.140)	(-8.145)	(-14.887)	(-8.061)
Category 5 * Realized	-56.429**	-0.262**	-55.594**	-0.260**	-55.117**	-0.258**
	(-19.783)	(-10.146)	(-19.326)	(-10.025)	(-19.134)	(-9.944)
Category 6 * Realized	-74.971**	-0.292**	-74.352**	-0.291**	-73.940**	-0.289**
	(-17.311)	(-9.422)	(-17.169)	(-9.372)	(-17.015)	(-9.299)
Category 7 * Realized	-86.682**	-0.242**	-85.762**	-0.239**	-85.284**	-0.238**
	(-11.178)	(-6.089)	(-11.128)	(-6.013)	(-11.140)	(-5.968)
Category 8 * Realized	-110.887**	-0.175**	-108.593**	-0.170**	-108.287**	-0.168**
	(-17.329)	(-6.264)	(-16.921)	(-5.599)	(-16.907)	(-5.560)
Population * Realized			-0.001**	0.000	-0.001**	0.000
			(-2.430)	(-1.038)	(-2.453)	(-1.062)
Infections * Realized					-1.323*	-0.003
					(-1.854)	(-0.709)
Deaths * Realized					0.005	0.000
					(0.219)	(-0.328)
Dependent Variable:	*PRICE*	*ln(PRICE)*	*PRICE*	*ln(PRICE)*	*PRICE*	*ln(PRICE)*
Num. Obs.	258792	258792	258792	258792	258792	258792
Adj/Psuedo R-squared	0.434	0.426	0.440	0.432	0.441	0.434

Logit estimates expressed as odds ratios. t-statistics (LPMs) or z-scores (logit) in parentheses.
** Significant at the 5% level. * Significant at the 10% level.

Table 6. Alternate Specifications

	(1)	(2)	(3)	(4)	(5)	(6)
Realized	1.201	-6.834**	-7.094**	-7.309**	1.540**	-6.492**
	(0.961)	(-3.027)	(-3.195)	(-3.304)	(2.106)	(-3.233)
Category 2 * Realized	1.695**	-10.404**	-10.487**	-10.453**	1.720**	-10.013**
	(2.097)	(-4.413)	(-4.516)	(-4.509)	(2.100)	(-4.867)
Category 3 * Realized	3.536**	-21.605**	-21.761**	-21.796**	3.640**	-20.430**
	(5.231)	(-9.330)	(-9.547)	(-9.580)	(5.197)	(-10.072)
Category 4 * Realized	3.445**	-35.962**	-35.997**	-36.255**	3.532**	-33.621**
	(5.362)	(-14.681)	(-14.887)	(-15.004)	(5.326)	(-16.401)
Category 5 * Realized	5.133**	-55.006**	-55.116**	-56.828**	4.686**	-50.409**
	(6.311)	(-18.978)	(-19.134)	(-19.404)	(5.857)	(-23.683)
Category 6 * Realized	12.373**	-73.898**	-73.944**	-77.183**	10.666**	-63.404**
	(8.023)	(-16.965)	(-17.015)	(-17.517)	(7.455)	(-25.752)
Category 7 * Realized	94.826**	-85.427**	-85.289**	-94.551**	74.577**	-67.628**
	(4.459)	(-11.197)	(-11.140)	(-11.800)	(4.215)	(-20.211)
Category 8 * Realized	27.760**	-108.452**	-108.254**	-107.900**	21.264**	-93.921**
	(4.837)	(-16.862)	(-16.878)	(-16.445)	(4.428)	(-12.825)
Population * Realized	1.000	-0.001**	-0.001**	-0.001**	1.000	-0.001**
	(-1.075)	(-2.420)	(-2.453)	(-2.306)	(-0.982)	(-8.540)
Infections * Realized			-1.323*	-1.449**	0.989	-1.005**
			(-1.854)	(-1.970)	(-0.239)	(-3.533)
Deaths * Realized			0.005	0.005	1.001	0.002
			(0.219)	(0.186)	(1.323)	(0.238)
Infection Rate*Realized	0.983	-0.721**				
	(-0.385)	(-2.673)				
Death Rate*Realized	1.001	0.003				
	(1.388)	(0.837)				
Advance*Realized			-0.367			
			(-1.209)			
Stay*Realized					0.288**	
					(-8.620)	
Model:	Logit	Linear	Linear	Linear	Logit	Selection
Dependent Variable:	*CLOSED*	*PRICE*	*PRICE*	*PRICE*	*CLOSED*	*PRICE*
Collapsed Data	Y	N	N	Y	Y	N
Num. Obs.	10506	258792	258792	8976	10506	286092
Adj R-squared	0.156	0.443	0.442	0.560	0.167	

Logit estimates expressed as odds ratios. t-statistics (price) or z-scores (logit) in parentheses.
** Significant at the 5% level. * Significant at the 10% level.

Table 7. Closure and Price Regressions, July

	(1)	(2)	(3)	(4)
Realized	0.006**	-0.049	-1.074**	1.176
	(2.835)	(-1.468)	(-3.193)	(1.174)
Category 2 * Realized		0.042		1.679
		(1.226)		(1.539)
Category 3 * Realized		0.050		1.960*
		(1.476)		(1.871)
Category 4 * Realized		0.063*		-0.430
		(1.846)		(-0.355)
Category 5 * Realized		0.065*		-7.506**
		(1.871)		(-4.803)
Category 6 * Realized		0.076**		-11.540**
		(2.110)		(-4.701)
Category 7 * Realized		0.049		-21.747**
		(1.434)		(-5.099)
Category 8 * Realized		0.048		-25.982**
		(0.830)		(-3.871)
Population * Realized		0.000		-0.001**
		(0.543)		(-2.860)
Infections * Realized		0.001		0.069
		(0.196)		(0.127)
Deaths * Realized		0.000		-0.005
		(-0.905)		(-0.240)
Model:	LPM	LPM	Linear	Linear
Dependent Variable:	*CLOSED*	*CLOSED*	*PRICE*	*PRICE*
Num. Obs.	10526	10506	300602	300602
Adj/Psuedo R-squared	0.000	0.008	0.000	0.453

t-statistics in parentheses. ** Significant at the 5% level. * Significant at the 10% level.