CESifo WORKING PAPERS

10849
2023

December 2023

Strategic Responses to Algorithmic Recommendations: Evidence from Hotel Pricing

Daniel Garcia, Juha Tolvanen, Alexander K. Wagner

CESifo

Impressum:

CESifo Working Papers
ISSN 2364-1428 (electronic version)
Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH
The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute
Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de
Editor: Clemens Fuest
https://www.cesifo.org/en/wp
An electronic version of the paper may be downloaded
· from the SSRN website: www.SSRN.com
· from the RePEc website: www.RePEc.org
· from the CESifo website: https://www.cesifo.org/en/wp

CESifo *Working Paper No. 10849*

Strategic Responses to Algorithmic Recommendations: Evidence from Hotel Pricing

Abstract

We study the interaction between algorithmic advice and human decisions using high-resolution hotel-room pricing data. We document that price setting frictions, arising from adjustment costs of human decision makers, induce a conflict of interest with the algorithmic advisor. A model of advice with costly price adjustments shows that, in equilibrium, algorithmic price recommendations are strategically biased and lead to suboptimal pricing by human decision makers. We quantify the losses from the strategic bias in recommendations using as structural model and estimate the potential benefits that would result from a shift to fully automated algorithmic pricing.

JEL-Codes: D220, D830, L130.

Keywords: advice, algorithmic recommendations, human decisions, adjustment cost, delegation.

Daniel Garcia
Department of Economics
University of Vienna / Austria
daniel.garcia@univie.ac.at

Juha Tolvanen
Department of Economics
Tor Vergata University Rome / Italy
juha.tolvanen@uniroma2.it

Alexander K. Wagner
Department of Economics
University of Salzburg / Austria
alexanderk.wagner@plus.ac.at

December 14, 2023
accepted for publication in *Management Science*
We are grateful to three anonymous referees, the department editor, and the associate editor for excellent comments. We also thank Emilio Calvano, Axel Ockenfels as well as participants at the 2021 CESifo Economics of Digitization Workshop, the EEA-ESEM 2022 and seminar participants at the TSE, CEU, UNLP, UPF, University of Zurich, University of Southampton, and University of Vienna for helpful comments and suggestions. The authors thank the anonymous corporate sponsor for access to the data used in this study. The computational results presented were obtained in part using the Vienna Scientific Cluster (VSC). Garcia gratefully acknowledges funding from FWF-FG6 Project. Wagner gratefully acknowledges financial support from the Anniversary Fund of the Oesterreichische Nationalbank (Project 18878).

1 Introduction

Organizations increasingly rely on prediction algorithms for decision making, with applications ranging from hiring policies at tech firms (Hoffman et al., 2018), bail decisions by judges (Berk, 2017; Kleinberg et al., 2018; Ludwig and Mullainathan, 2021), to pricing (Chen et al., 2016b; Miklós-Thal and Tucker, 2019; Calvano et al., 2020; Garcia et al., 2022; Hansen et al., 2021; Harrington Jr, 2022; Brown and MacKay, 2023; Abada and Lambin, 2023; Assad et al., forthcoming). In many of these applications, machines augment human decisions by acting as advisors to human managers who retain the final decision right. It is well understood that advice can be successful only to the extent that incentives of the human decision maker and the advisor are sufficiently aligned.[1] Although recent theoretical work has considered the effects of misaligned incentives in human-algorithm interactions on communication (Cowgill and Stevenson, 2020) and delegation decisions (Athey et al., 2020), empirical evidence on how algorithmic advice and human decision making work together in strategic situations remains scarce at best.

In this paper, we study the interaction between algorithmic price recommendations from an independent advisor and actual price setting behavior by human managers in a hotel-room pricing context. We develop a parsimonious model of advice that allows us to identify a novel source of conflict between the algorithmic advisor and the human decision maker originating from price-adjustment costs humans face. This *strategic* conflict of interest can lead to substantially distorted recommendations by the algorithm and suboptimal pricing decisions by human managers. We believe this is the first paper to model and estimate a conflict of interest in communication arising from differences in adjustment costs. This conflict of interest is likely to be ubiquitous, especially in algorithmic recommendation systems where the 'human in the loop' has access to an easily implementable default option (see section 8 for examples).

For our empirical analysis, we leverage a dataset containing millions of algorithmic price recommendations of an independent revenue management company (the algorithm's designer), prices set by hotel managers (human manager), and the corresponding bookings from nine different hotels. Using this data, we estimate a structural model of advice and quantify the resulting losses from mispricing. Our main counterfactual experiment shows that full delegation to the recommendation algorithm significantly outperforms the current

[1] See Sobel (2013) for an excellent summary of the literature on advice with an informed sender and uninformed receivers. We discuss the contributions of our study to this literature in section 2.

organizational setting in which human managers set prices. The status quo setting reduces the expected loss from mispricing by 1 to 2 percent compared to not changing prices at all, while delegation to the algorithm would reduce the expected loss across hotels by 4 to 36 percent.

In our setting, both the revenue management company and the hotel manager benefit from maximizing the hotel's profits.[2] The primary source of conflict we consider is not due to ex-ante misaligned objectives between the two agents, but stems from the fact that the recommendation algorithm faces no adjustment costs when changing prices, whereas the hotel manager does. The algorithm would therefore like to implement price updates more frequently than what is privately optimal for the hotel manager.[3] The algorithm's designer has an incentive to offer recommendations that *exaggerate* the change in the optimal price to induce the human manager to update prices faster, because larger deviations from the optimal price result in larger expected revenue losses. This strategic bias, in turn, has two effects. On the one hand, both agents incur a welfare loss whenever a more biased recommendation is accepted by the manager. On the other hand, the designer of the algorithm benefits (more than the manager) whenever exaggerated recommendations prompt more frequent price updates from the manager.

We build a stylized model of advice that captures the key features of the price-setting interaction of the algorithmic advisor and the manager observed in the data. In the model, we assume that human managers incur price-adjustment costs. The manager costlessly observes the price and the algorithmic recommendation of a given product, defined as a specific room-arrival-date combination. She then decides whether to keep the current price at no additional cost or allocate costly attention to the pricing task. If she chooses the latter, she receives an informative signal about the optimal price and then decides whether to copy the recommended price. If she does not copy the recommendation, she learns additional information by incurring cognitive 'thinking' costs and then gets to update prices freely.[4] The assumed timing and cost structure for adjusting prices is also reflected by the pricing interface managers use: accepting

[2] This shared objective is motivated by the fact that the revenue management firm compares their customers' revenues before and after they started using the recommendation algorithm, and is heavily marketing these benchmarks to new and existing customers. Consequently, higher induced revenues help the revenue management firm to attract new customers and retain existing ones.

[3] The difference in how these two experience adjustment costs is evident from our descriptive analysis in section 3, showing that human managers update prices much less frequently than recommended by the pricing algorithm.

[4] Our price-adjustment model is consistent with different interpretations, including menu costs and managerial information-processing costs, as we discuss in section 5. There is also a substantial macroeconomic literature studying adjustment and information processing costs in price setting, see e.g., Alvarez et al. (2011).

a set of price recommendations is relatively easy and requires only a single click whereas adjusting prices freely entails accessing different screens and imputing prices manually.

Our model of strategic advice with costly price adjustments successfully matches the key empirical relationships between algorithmic recommendations and realized prices summarized in section 4. First, managers' price updates are much less frequent than updates in price recommendations. Over the booking horizon, human managers update prices for a particular product about once a month whereas algorithmic recommendations are updated four times more often. This difference in updating frequencies, together with rarely observed small price changes by managers, indicate that they are facing considerable price-adjustment costs.

Second, our model shows that, for each realization of the price recommendation, there exists a cutoff value such that the manager devotes attention to pricing only if her adjustment cost shock is below the cutoff. This cutoff is increasing in the size of the recommendation change. The monotonicity induces a positive correlation between adjustment costs and the magnitude of the recommendation change, conditional on the manager changing the price. It follows directly that, in case of updating prices, the manager is more likely to face higher adjustment costs the larger the change in the recommendation. If the two adjustment costs are correlated, she is more likely to copy the recommended price rather than adjusting prices manually.[5] This is consistent with the empirical observation that, conditional on changing prices, copying the recommendation becomes more likely than manual adjustments the larger the change in the recommendation.[6] This pattern is otherwise hard to reconcile with standard models of advice as we discuss in section 4 (see also footnote 16 and appendix H).

Our third, and most salient reduced-form observation is that managers only partially incorporate the information contained in the recommendation. When the manager updates prices manually, a one-percent change in either direction in the recommended prices leads, on average, to a 0.72 percent change in the same direction in the realized price. In line with this empirical finding, the pass-through of recommendations into actual prices is imperfect in our model because the human manager expects a biased price recommendation. The strategically biased recommendation makes it more profitable for the manager to manually update prices whenever the direction of the price change suggested by her private information contradicts the one recommended by the algorithm. This *negative selection* of hotel manager's private

[5] Correlation between the two adjustment costs is highly plausible because both represent an opportunity cost of the manager's time or cognitive resources.

[6] For example, the average probability of copying the recommendation conditional on a price change is 84% and increases to 95% once the change in the recommendation exceeds 10% of the current price.

information in case of manual price updates further decreases the empirically observed pass-through rate of recommended prices.[7] The reason is that manual price changes are more profitable in situations in which the manager's information strongly contradicts the algorithmic advisor's signal. Taken together, the assumed timing and the two adjustment costs are required for a model rich enough to generate the counterintuitive empirical pattern of (i) a high rate of copying recommendations as well as (ii) a considerably dampened pass-through of algorithmic recommendations in case the manager decides to adjust prices manually.

Two key ingredients of our model presented in section 5 are the biased recommendations of the algorithm and how they are perceived by managers. We posit that the algorithm aims to induce revenue-maximizing decisions and holds correct expectations about the manager's response to different recommendations. The algorithm's designer chooses a reporting strategy that, for tractability, we assume to be a linear factor, multiplying the change in the privately observed component of the optimal price. If this factor exceeds one, the algorithm exaggerates its private information. In equilibrium, the hotel manager is assumed to have correct beliefs about the bias factor, and therefore forms accurate expectations about the information held by the algorithm.

In section 6, we estimate the model parameters using a minimum distance estimator while requiring that there are no beneficial deviations for the algorithm's designer from the chosen bias factor. For the complete sample of hotels in our data, the bias factor we identify is 1.2. The bias in recommendations is significantly larger than 1 but lower than the naive estimate of 1.39 necessary to explain the low pass-through of 0.72% in a model without selection. Estimating the model at the hotel level, we observe a considerable degree of heterogeneity across hotels, with estimates of the bias ranging from 1.05 to 1.5. In addition, we find that hotel managers have access to potentially very precise information but only rarely make use of it, implying that they must be facing substantial adjustment costs. This leads to inaccurate decisions exemplified by the fact that the standard deviation of the actual prices is smaller than the standard deviation of the difference between the estimated optimal and actual prices. In other words, most of the variation in the optimal price is not captured by the price updates implemented by the manager.

We then take our estimated model to study whether full delegation to the algorithm would lead to better pricing decisions in section 7. Delegation has obvious benefits as it brings about instantaneous and costless decision making and aligns preferences of the algorithmic advisor

[7] In other words, our model predicts that the actual bias in recommendations is smaller than the gap between recommended changes and manual price changes.

and the human manager. However, the hotel manager's private information is ignored as an input for decision making when decisions are delegated. Our counterfactual experiments show that the status quo of no delegation is only slightly better than complete inaction (no price change at all), while delegating pricing to the algorithm would reduce expected losses by up to 36 percent in the sample. There is clear value in delegating to the algorithm, as it provides a cheap way of adjusting prices more frequently and incorporating additional price information. Decomposing these gains shows that the majority of gains from delegation to the algorithm stems from increased flexibility due to more frequent pricing decisions (80%). The remaining gains come from de-biasing recommendations (10%) as well as costless information processing (10%). Finally, we also vary the relative price information held by the algorithm and the manager. The counterfactual analysis reveals that the manager becomes better off in the status quo when the algorithm's relative share of information increases. But delegation would still be optimal for almost all levels of relative information held by the algorithm due to the large costs managers face for adjusting prices.

The paper proceeds as follows. Section 2 relates our contribution to the existent theoretical and empirical literature on (algorithmic) advice in strategic settings. In section 3, we provide important details of our pricing data as well as background information about the decision making environment, before presenting key empirical facts about algorithmic price recommendations and actual prices chosen by human managers in section 4. Section 5 develops a model of information processing consistent with the main stylized facts. Estimation results of the model are presented in section 6. Section 7 provides results of counterfactual experiments when pricing it delegated fully to the recommendation algorithm. Section 8 concludes with a discussion of our main findings and comments on the general difficulty of solving frictions that arise in strategic interactions between a (third-party) algorithmic advisor and a human decision maker.

2 Related Literature

This paper adds to a thriving literature studying the interaction between algorithmic advice and human decisions in an economically important application of managerial pricing. We develop and empirically estimate a model of strategic communication between an algorithmic advisor and a human decision maker. More generally, the present paper is one of the only empirical papers estimating a cheap-talk model with observational data.

There exists a large theoretical literature on strategic advice in organizations (e.g. Kamenica and Gentzkow, 2011; Sobel, 2013; Kamenica, 2019).[8] In the canonical setup (see Crawford and Sobel, 1982) an informed agent communicates with an uninformed principal who has to make a decision that affects both the agent and the principal's payoffs. The agent and the principal have only partially aligned interests because they disagree on the (ex-post) optimal action. Instead, we focus on the principal's adjustment cost as the source of disagreement. We believe this is a relevant consideration in many organizations, whereby generalist managers rely on information from several experts. In the context of the interaction between human decision makers and machine advisors, this is almost always guaranteed to occur, as adjustments costs of machines are infinitesimal compared with those of humans. In the spirit of Aghion and Tirole (1997), machines hold *real authority* because the information-processing costs of the human decision maker vastly exceed those of the algorithmic advisor. A similar tension arises in the cheap talk model of Kartik et al. (2007), in which a fraction of the audience is *naive* and takes the message at face value. In equilibrium, senders exaggerate their claims so that the marginal incentive to misrepresent their information to sophisticated receivers equals the cost they bear on the naive ones. In our setting, all receivers are sophisticated but behave naively to economize adjustment costs.

To the best of our knowledge, there exist only two papers that explicitly incorporate inattentive decision makers in a model of advice. Agrawal et al. (2019) studies the impact of artificial intelligence on human decision making. In their setting, a principal has to choose whether to implement a new project with uncertain costs and benefits and has access to truthful information provided by a machine (à la Aghion and Tirole, 1997). In the presence of a large number of different projects, the human decision maker tends to focus her attention on high-stake projects and fully delegates decision making to the machine in those with low stakes. In our setting, we observe hotel managers choosing a price which exactly matches the recommendation (akin to delegation, or rubber-stamping) *more* often, whenever the recommended price is further from the current price. The crucial insight is that rubber-stamping also requires, in contrast to fully automated decision making, at least some attention from the human principal.[9]

Bloedel and Segal (2020) study a persuasion model where the principal is subject to

[8] Sobel (2013) provides a survey of cheap talk communication and Kamenica (2019) gives an overview of recent advances in Bayesian persuasion and information design, where the sender can commit in advance to an information structure. Our model lies somewhere in-between, as the agent chooses a linear reporting strategy but can secretly deviate from it.

[9] In the algorithmic pricing industry, some companies offer an arrangement similar to the one suggested in Agrawal et al. (2019); that is, the algorithm directly implements changes if they fall within a given (target) price range, while human approval is needed if the suggested price falls outside the defined range.

rational inattention.[10] As in our setting, inattention induces a moral hazard problem that leads the advisor to distort her messages to motivate the principal to pay attention. They show that full disclosure is optimal only if stakes are low, and instead pool medium and high stakes. In the present study, we consider only linear reporting strategies and leave the full-fledged analysis of the sender for future work. In any case, we do not observe pooling or bunching. On average, prices set manually by the hotel manager increase continuously in the recommendation in our data.

More broadly, we contribute to the empirical literature on strategic communication in organizations.[11] We are aware of two papers that use equilibrium analysis to identify strategic communication behavior. Backus et al. (2019) provide evidence of strategic communication in bargaining in a large online marketplace in which impatient sellers use round numbers in their posted price as a signaling device. Camara and Dupuis (2014) study movie reviews through the lens of a reputational cheap-talk model, uncovering a significant conservative bias. Our setting has several advantages. First, both the sender and the receiver are professionals and face serious financial consequences from their actions. Second, there is an obvious mapping between messages and recommended actions in our data. Third, the action space is sufficiently rich to directly identify the posterior beliefs of the receiver whenever she chooses a price that departs from the recommendation.

Finally, we also contribute to the literature on algorithmic bias and human decision making. Most of these papers consider algorithmic predictions as potential substitutes of human experts, assessing their potential advantages (accuracy, speed) and disadvantages (algorithmic bias or negative perception of third-parties).[12] Two notable exceptions are Bundorf et al. (2019) and Caro and Sáez de Tejada Cuenca (2023). The first paper studies the impact of algorithmic recommendations on the purchasing decision of health insurance plans among the elderly. As in our setup, they find human inertia to be a major concern, but their algorithmic recommendation is assumed to be non-strategic. The second paper provides reduced-form evidence about the way sales managers at a large fashion retailer react to algorithmic price recommendations during sales campaigns. They show that cognitive workload related to the

[10] There is a large literature on rational inattention, with some important applications to organizations (see Maćkowiak et al., 2023, section 3.3).

[11] There is a growing literature that empirically studies persuasion, from advertising to mass media, see DellaVigna and Gentzkow (2010) for an excellent summary. Most of these papers focus on identifying the persuasion effect (or lift ratios). Instead, we attempt to uncover the economic incentives underlying persuasion and explore counterfactual arrangements that may improve decision making.

[12] Prominent examples in this stream of literature are Hoffman et al. (2018), Kleinberg et al. (2018), Ribers and Ullrich (2023), Chan et al. (2022) and Currie and MacLeod (2017).

complexity of the task (the number of prices to be set) is a key driver of adherence to algorithmic recommendations, much like what we see in our data. Furthermore, similar to the first paper, they do not consider any strategic interaction between the algorithm's designer and the manager. Our study strongly suggests that assuming truthful recommendations as a counterfactual scenario may be neither optimal nor realistic.

3 Data and Setting

The data for our analysis contains almost 6 million observations of hotel-room pricing information, including algorithmic recommendations, actual prices set by human decision makers and the corresponding universe of about 60 thousand bookings, all aggregated at the daily level. This high-resolution, proprietary data is provided by an anonymous corporate sponsor, who is based in Europe and provides revenue management services to hundreds of independent hotels. The pricing and booking data come from 9 different hotels, eight of which are located at resort destinations (hotels A to H) and one in an urban area (the hotel I). These hotels were selected because (i) they have substantial experience with the recommendation system and (ii) they are *not* located next to another hotel who is a client of the revenue management company, reducing concerns about algorithmic collusion. Our booking and pricing data spans for each hotel over a period of about 14 months. We observe for each room and each possible arrival day the flow of bookings, the recommended price by the revenue management service and the actual price charged by the hotel. The actual price is an index price which determines, together with possibly channel-specific discounts or surcharges, the final price of a room. The revenue management system's algorithm and the hotel manager rely on this price as the main instrument for price optimization. See Garcia et al. (2022) for more details on the data and institutional background.

A key input for our analysis is the algorithmic price recommendation. The recommendation algorithm is provided by the revenue management service and aims at maximizing hotel revenue through optimized pricing. The hotels pay the revenue manager a fixed fee but the revenue management firm heavily uses its success in increasing its customers' revenues when it markets its services to both new and existing customers. Hence, the firm is highly motivated to increase its customers' revenues. The algorithm uses all booking information and collects additional demand-related information including, among others, local variation in weather conditions, events, public holidays, hotel reputation, and competitor prices. In addition, the revenue management service and the hotel manager exchange information about local demand conditions regularly. The algorithm processes all available information and then generates a price recommendation for each product, i.e. room-arrival-date combination.[13]

The hotel manager decides every day whether to use the revenue management system to update prices. If she logs into the system, the dashboard displays for each room the current recommended price and the actual price. She then decides which price to update and by how much. Although the hotels in our data are representative of their respective regions, with about 50 rooms each, they are relatively small by international standards. According to private communication with the revenue management service, hotels are family-run and thus managing prices takes only a small fraction of a hotel manager's weekly workload. One of the main selling points of the recommendation service is to simplify and reduce this workload. Appendix C reproduces some of the evidence from Garcia et al. (2022) on the opportunity cost of adjusting prices faced by the hotel managers and on how accepting recommendations consumes less time than adjusting prices manually.[14]

Our analysis relies on recommendation and price changes as the main variables, see Table 1 for descriptive statistics. From the panel of daily prices, we construct the first differences in the log price and define an *update* whenever this difference is non-zero. We define the change in the recommended price as the change in the logarithm of the algorithmic price recommendation since the last price update. We restrict the full sample to include only observations for which the initial price matched the recommended price for our analysis ("Main Sample" in Table 1). This selection allows us to interpret differences between the price and the recommendation as differences in the current information processing and removes any feedback effects from past prices into future price recommendations. The resulting main sample includes approximately 34% of observations and 58% of the price updates of the full sample.[15]

4 Stylized Facts: Recommendations and Prices

In this section, we present key empirical facts about the relationship between price recommendations by the algorithm and price updates by the human manager. These observations

[13] Although we do not have access to the proprietary pricing algorithm, it is sufficient for our empirical analysis that the recommended price of the algorithm contains some relevant information for the hotel manager.

[14] Similarly, Huang (2022) argues that Airbnb hosts face significant adjustment costs when adjusting their prices. Although our family hotel managers are perhaps more professional and have access to slightly better sources of price information, their task is also more complex because a hotel typically has multiple different room types. Still, we think that the two settings are fairly comparable in terms of price-adjustment costs.

[15] We also provide evidence that the qualitative features of the pass-through rate in the main sample shown in Table 4 also hold for the full sample in appendix A. If anything, the magnitude of the implied strategic bias in recommendations would be larger in the unrestricted sample.

Table 1: Price Updates and Recommendations

	Main Sample			*Full Sample*		
	Min	Mean	Max	Min	Mean	Max
Update Rate	0.012	0.038	0.045	0.005	0.022	0.048
Update Rate | Rec No Change	0.001	0.006	0.009	0.004	0.016	0.025
Update Rate | Large Rec	0.018	0.143	0.658	0.006	0.046	0.676
Update Copy Rec	0.757	0.840	0.894	0.497	0.665	0.925
Update Copy Rec | Large Rec	0.869	0.950	1.000	0.431	0.860	0.931
Update Size	0.033	0.048	0.066	0.035	0.055	0.228
Update Size | Copy Rec	0.029	0.044	0.061	0.060	0.143	0.208
N	2,017,932	2,017,932	2,017,932	5,916,580	5,916,580	5,916,580

Notes: For all statistics we report the maximum, minimum and average value across hotels for the main sample and the full sample. Update Rate is the proportion of products in which we observe a price update on a given day. We report in rows 1 to 3 the update rate unconditionally, conditional on the recommendation not having changed (Rec No Change) and conditional on the recommendation having changed by at least 10% (Large Rec). Update Copy Rec rate is the proportion of updates in which the updated price matches the recommendation exactly. We report in rows 4 to 5 the Update Copy Rec unconditionally and conditional on an absolute change in the recommendation of at least 10% (Large Rec). The Update Size is the average log change in the realized price following an update. We report in rows 6 to 7 the Update Size unconditionally and conditional on matching the recommendation exactly (Copy Rec).

inform the choice of our model of price adjustments we present in section 5. For the following descriptive analysis, we pool observations across all hotels in our sample, but the findings also hold qualitatively for each hotel individually.

Observation 1. *Price updates are much less frequent than updates in price recommendations.*

Managers adjust prices only infrequently, on average, once every 35 days, with considerable heterogeneity across hotels as shown in Table 1. Because algorithmic price recommendations change much more frequently, once every seven days, the difference in updating frequencies leads inevitably to a divergence between recommendations and actual prices over time. The inertia in updating prices is also reflected in the distribution of price changes, shown in Figure 1, and exhibits little mass around 0. This pattern is a first indication that price-setting human managers face considerable adjustment costs (Nakamura and Steinsson, 2008).

Observation 2. *The frequency of a price update is positively related to the size of the recommended price change.*

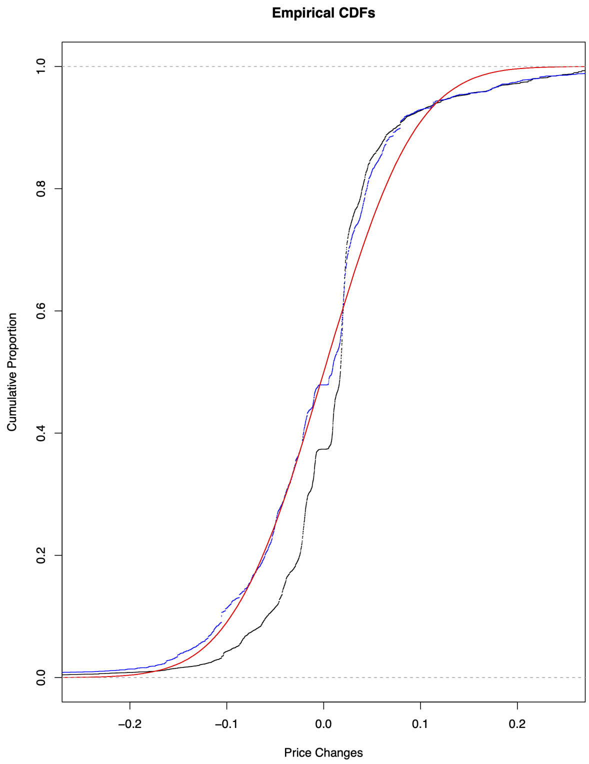

Figure 1: Distribution of price updates (in log changes)

Notes: The black line plots the empirical cdf of price changes. The blue line depicts the empirical cdf of manual price changes. The red line plots a normal cdf with the same standard deviation as the black distribution.

Larger changes in the recommendation are associated with a higher likelihood of a price update, as shown in Figure 2. For instance, if the recommended price has remained unchanged since the last price update of the hotel manager, the probability of a price update today is less than 1%. However, if the current recommendation is outside a ten-percent band of the original recommendation, the probability of a price update exceeds 11%. This positive correlation is also confirmed by fixed-effects regressions in Table 2 which account for variation across different products of the same room type (e.g. standard room), hotel, and arrival-month for a given date.

Observation 3. *The probability that a price update copies the recommended price is increasing in the size of the recommended price change.*

Conditional on observing a price update, hotel managers are very likely to update the price to exactly match the price recommendation. On average, around 85% of the price up-

Table 2: Price Update Probability

	Update Probability			
Rec Change	0.033***	0.105***	0.109***	0.108***
	(0.004)	(0.003)	(0.003)	(0.003)
Rec Update	0.110***	0.123***	0.128***	0.128***
	(0.000)	(0.000)	(0.000)	(0.000)
Hotel × Date FE	No	Yes	No	No
Room × Date FE	No	No	Yes	No
Room × Date × Month FE	No	No	No	Yes
N	2,017,929	2,017,929	2,017,929	2,017,929

Notes: Fixed-effects regressions. The dependent variable is the instantaneous probability of a price update. Rec Change is the cumulative (log) change in the recommendation since the last price update. Rec Update is a dummy which takes the value one if the recommendation has changed since the last price update. Room is the room type, Date is the booking date and Month refers to the arrival month. Significance levels: *** $p < 0.001$

dates copy the currently recommended price perfectly; with some heterogeneity across hotels. The probability of copying the recommended price is even higher if one conditions on a large change in the recommended price as can be seen in Figure 3. In particular, if the recommendation change exceeds 10%, the hotel manager chooses a price that exactly matches the recommendation with 95% probability. Table 3 shows that this positive correlation also remains in a fixed-effects regression that exploits only variation across neighboring arrival dates for the same booking date. Importantly, the updating pattern of hotel managers, summarized in Observation 2 and 3, is inconsistent with standard models of advice in which the influence of the (algorithmic) advisor decreases when making more *extreme* recommendations.[16] Together with the other empirical facts, we will account for this distinctive updating behavior

[16] The feature of decreasing effectiveness of more extreme advice comes in various forms in the literature. In cheap-talk games, it results in less precise communication. The same comparative static holds in games that introduce reputational, moral, or strategic concerns of lying (Kartik et al., 2007). Similarly, the principal rubber-stamps decisions in Aghion and Tirole (1997) under contingent delegation that involve low stakes but assumes control when stakes are high, thus reducing the influence of the agent. Finally a robust finding in the newsvendor literature is that decision makers under-react to large shocks (pull-to-center effect) when choosing capacity (Schweitzer and Cachon, 2000; Bolton and Katok, 2008; Bolton et al., 2012). In summary, all of these models are inconsistent with our observation that managers are more likely to copy price recommendations the more recommendations deviate from the current price.

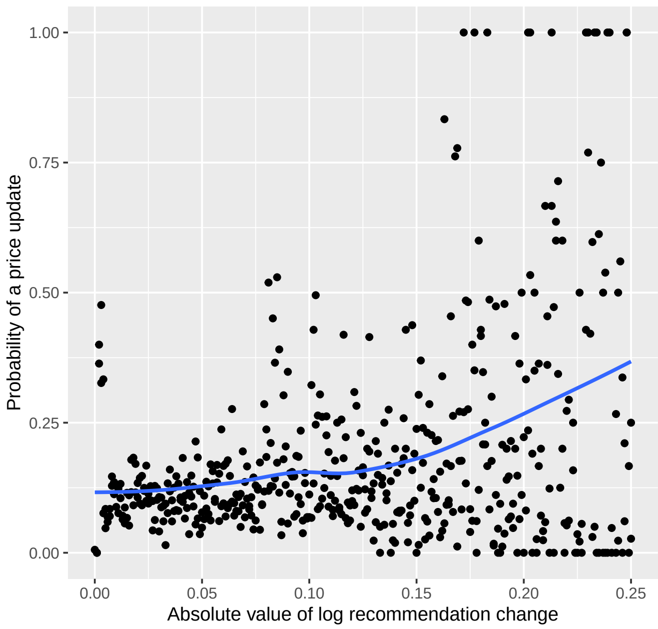

Figure 2: Update Rate

Notes: Each point represents a 0.001-sized bin. The horizontal axis captures the log change in the recommendation. The vertical axis contains the average probability for that bin.

of managers in our model in section 5.

Observation 4. *There is only a partial pass-through of the change in the recommendation into actual prices.*

If the interests of the hotel manager and the recommendation algorithm were perfectly aligned and the manager's arrival of private information would be uncorrelated with the direction of her private information, one would expect that, on average, a one Euro increase (decrease) in the recommendation would bring about a one Euro increase (decrease) in the price. The observed difference between the recommendation and the actual price could then be attributed to the additional, idiosyncratic information held by the manager. Instead, we observe as shown in Table 4 a much lower pass-through rate of 72.5%. In other words, when hotel managers manually update their prices, the average price change only partially reacts to the change in the recommended price. Including various controls, such as room type-arrival week fixed effects and a polynomial of the days before arrival, leads only to a modest increase in the estimated coefficient (73%). It follows that hotel managers must believe that the pricing algorithm exaggerates the optimal price change on average.

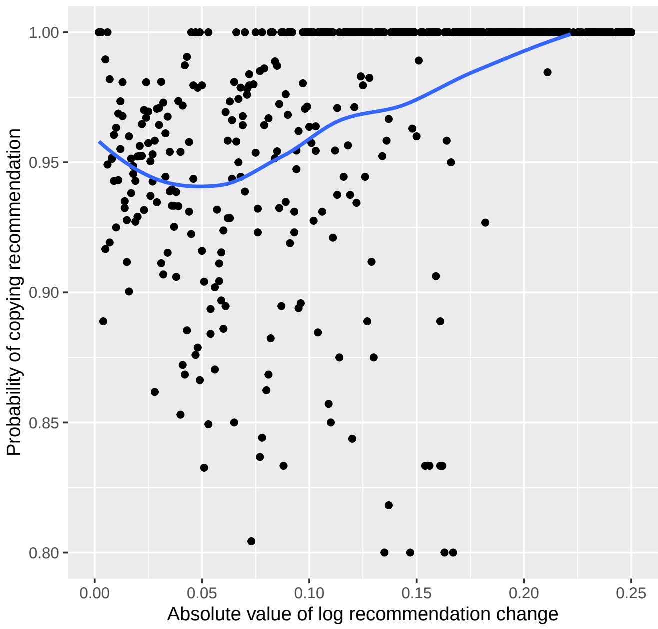

Figure 3: Matching the Recommendation

Notes: Each point represents a 0.001-sized bin. The horizontal axis captures the log change in the recommendation. The vertical axis represents the proportion of updates that exactly match the recommendation.

Interestingly, the unconditional relation between recommended prices and actual prices is continuous and almost linear, see Figure 5. This fact is inconsistent with equilibria in standard cheap-talk models, which display discontinuities (bunching) to ensure incentive compatibility. It is also at odds with multi-dimensional models of communication in which the size of the recommendation change *signals* the quality (precision) of the information held by the advisor, thus inducing a higher likelihood of copying when the recommendation is further from the current price. As the size of the recommendation change increases, the marginal impact on the posterior belief of the manager should increase, regardless of whether the manager copies it.

5 Model

In this section we introduce a model of price adjustments with algorithmic recommendations and costly information acquisition of the hotel manager. The price-adjustment model rationalizes the empirical facts about the relationship between recommendations and pricing decisions presented in the section 4. To ease the mapping of the model to the data, we

Table 3: Price Update Copy Rates

	Copying Probability			
Rec Change	0.133***	0.141***	0.151***	0.151***
	(0.012)	(0.011)	(0.011)	(0.011)
Hotel × Date FE	No	Yes	No	No
Room × Date FE	No	No	Yes	No
Room × Date × Month FE	No	No	No	Yes
N	76,090	76,090	76,090	76,090

Notes: Fixed-effects regressions. The dependent variable is the probability of copying the recommended price. Data is restricted to neighboring arrival dates for a given booking day. Rec Change is the cumulative (log) change in the recommendation since the last update. Room is the room type, Date is the booking date and Month refers to the arrival month. Significance levels: *** $p < 0.001$

normalize all variables to refer to percentage changes since the last update.

5.1 Model Description

We begin by introducing the main elements of the model. A hotel managers intends to maximize profits, defined as $\Pi = \Pi_0 - \eta(p - p^*)^2$, where p is the current price, p^* is the optimal price given demand and cost conditions, and $\eta > 0$ is a parameter.[17] The optimal price is determined by $p^* = x + y + z$. Random variables x, y, and z are drawn independently from a symmetric distribution with zero mean and variance σ_i^2, for $i = x, y, z$.[18] The model assumes normal distributions for x, y, and z, although the main argument does not depend on the distributional assumption.

Figure 4 illustrates the hotel manager's sequential information-acquisition process and pricing decision in the task. Once the manager accesses the pricing interface, she learns the current price, normalized to $p = 0$, the recommendation r and the realized costs c_1 and c_2 for learning information y and z. Information x is the algorithm's private information regarding the optimal price and is not known to the hotel manager. The manager only learns about x

[17] This profit function can be micro-founded assuming a log-linear demand and semi-elasticity η. However, the model is more general than this, as we remain agnostic on whether this price maximizes revenue, profits, reputation or even some dynamic aggregate of the three.

[18] To see that independence is not a severe restriction, suppose that the manager observes $y = \alpha x + (1-\alpha)\tilde{y}$ so $corr(y, x) = \alpha$. Because the manager observes r, she can infer x, and therefore can easily extract $\tilde{y}$ which is now, by definition, independent of x.

Table 4: Pass-Through Rates of Recommendation

	Change in actual price			
	All	Manually Updated		
Rec Change	0.974***	0.725***	0.733***	0.738***
	(0.002)	(0.005)	(0.006)	(0.006)
Days ahead Polynomial	No	No	Yes	Yes
Room × Month FE	No	No	No	Yes
N	76,090	76,090	76,090	76,090

Notes: Linear regression model. The dependent variable is the cumulative change in the actual price since the last price update. Recommendation is the cumulative (log) change in the recommendation since the last price update. All regressions include all updates. Coefficients for *Manually Updated* correspond to the interaction term of recommendation × manual. *Room* is the room type and *Month* refers to the arrival month. Significance levels: *** $p < 0.001$

by observing the algorithm's price recommendation $r(x)$.

The manager then decides whether to allocate attention to adjusting prices. In case she does not, the current price, $p = 0$, is maintained and the manager incurs no cost. In case she allocates attention to the pricing task, she learns information y for the *attention cost* c_1. We think of y as information about a particular product which directly springs to the manager's mind, e.g. the hotel's chef is on a leave of absence. After the manager learned y, she can either choose to update the price by copying the recommendation resulting in $p = r$, or to acquire additional information z for the cognitive *thinking cost* c_2. Only in the case of learning z for cost c_2, the manager can update the price freely, such that, $p = E(p^* \mid r, y, z) = E(x \mid r) + y + z$.[19]

For parsimony, we assume that costs c_1 and c_2 are determined by a common cost shock c, drawn from a distribution $F(c)$. In particular, we assume that $c_i = b_i c$, with $b_i > 0$.

[19] Although the model seems rather intricate, it allows two simpler special cases. First, if $\sigma_y = 0$, the hotel manager has access to only a single costly signal. The manager has to decide upon observing the recommendation whether to do nothing, copy it at cost c_1 or incur both c_1 and c_2 to obtain further information and change the price manually. Second, if $c_1 > 0$ but $c_2 = 0$, the manager incurs only an attention cost and always acquires all the available information. In this case the manager would never copy the recommendation but would either stay inactive or manually adjust the price. However, these two hypotheses are rejected in our empirical results and hence we conclude that these simpler alternatives are unlikely to represent accurately the hotel managers' pricing problem. For a discussion of other alternative modelling assumptions, see section H in the appendix.

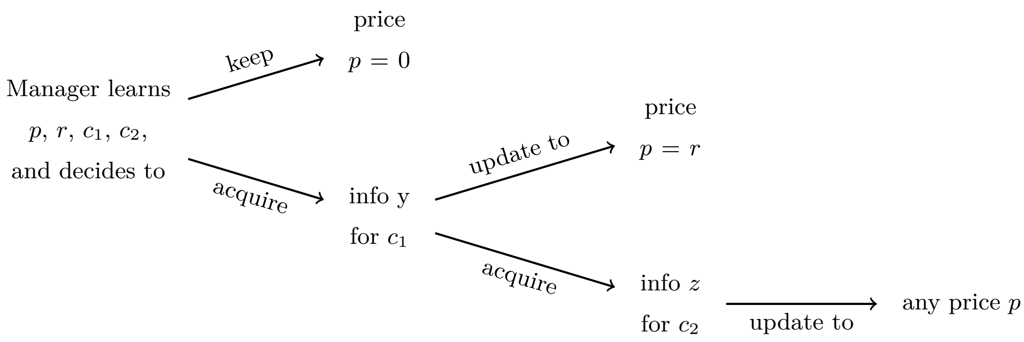

Figure 4: Information acquisition and pricing of manager

We think of c as cognitive costs with b_i measuring the difficulty of the task. Although the interpretation of the two adjustment costs, attention cost and thinking costs, is intuitively appealing we remain agnostic about the exact psychological nature of the two costs because our data does not allow to discriminate between different interpretations. An alternative interpretation of the cost structure would be to understand parameter c as the opportunity cost of a unit of time for the manager and b_i as the time required to learn the information.[20]

Under both interpretations, it seems easier for a manager to decide whether the recommendation is satisfactory than to fully determine the optimal price manually.[21] This structure is reinforced by the pricing interface, which allows accepting recommendations with a single click, while freely adjusting prices requires the manager to access an additional screen and enter each price manually. Neither our theoretical nor empirical model imposes however an assumption on which of the two adjustment cost is higher. Note also that the assumed sequential structure of the adjustment cost model with two signals and costs respectively is required for matching the empirically observed price updating pattern presented in section 4. We show in appendix D that a one-step structure of adjustment costs, e.g. without attention cost c_1, would be inconsistent with the observed price updating pattern in the data.

Another key ingredient of the model is the algorithmic price recommendation. As discussed earlier, we assume that the algorithm's designer cares about the hotel's profits directly,

[20] Another interpretation of the adjustment costs is that they may originate from the belief that consumers may react negatively to price variation, see Rotemberg (2005). If this was indeed the origin of adjustment costs, they should be directly incorporated into profits and the sluggishness of price adjustments may be profit-maximizing. Consumers, however, do not know whether price changes are the result of a change in the recommendation (c_1) or a manual adjustment (c_2). Thus, if the estimated cost component of the manual price adjustment (b_2) is significantly larger than that of copying (b_1), it is reasonable to conclude that the bulk of costs is due to cognitive thinking costs of managers –rather than an attempt to please consumers.

[21] This is also consistent with decision models under limited attention (e.g. Dean et al., 2017), showing that status quo bias is more likely in larger choice sets (in our case if the manager can choose between copying the recommendation and manually setting a price p).

$\Pi = \Pi_0 - \eta(p - p^*)^2$. The crucial difference between the algorithm's designer and the hotel manager is that the algorithm does not face any adjustment costs because recommendations are fully automatized. This difference creates a strategic conflict of interest between the designer and the manager: the designer would like the hotel manager to update more frequently than what is optimal to the manager. The hotel manager's update frequency can be influenced by strategically choosing recommendations $r(x)$. We analyze the perfect Bayesian equilibria of this game, restricting the algorithm's message space to linear functions of its posterior belief about the optimal price, $r = \frac{1}{\lambda}x$ with the bias factor of the recommendation $\lambda > 0$.[22]

Although the hotel manager does not directly observe λ, in any perfect Bayesian equilibrium she will form correct expectations about it. Notice that a further exaggeration from any given λ has two effects. First, given the hotel manager's (equilibrium) beliefs, she will incorrectly think that the optimal price has changed more than it really has, which induces a higher probability of the hotel manager changing the price manually. For small deviations this benefits the algorithm's designer. However, the hotel manager may also copy the recommended price. The larger is the total exaggeration, the larger is the hotel manager's pricing mistake in this case. In equilibrium, λ exactly equates this trade-off between more frequent updates and larger mistakes when copying recommendations at the margin.

We decided not to include additional biases in the objective function of the algorithm. The reasons for this are twofold. First, we want to focus the analysis on adjustment costs as they are the main novelty of the present study. Second, we do not find any tendency of the algorithmic recommendations to be biased either upwards or downwards which could point at misaligned objectives between the algorithmic advisor and the hotel manager: recommendations are, on average, approximately equal to the realized prices.[23] If the algorithm cared about revenue rather than profits, as in the example of pricing by Airbnb, see Huang (2022), we would see a tendency to recommend lower price levels. If, instead, the algorithm wanted to induce the manager to collude with other hotels as in, for example, Calvano et al. (2020) and Miklós-Thal and Tucker (2019), we would observe higher price recommendations by the

[22] This restriction on the algorithm's strategy space makes both the theoretical model and the empirical model more tractable. Furthermore, we illustrate in appendix E that the loss from restricting oneself to a linear reporting strategy is negligible for the algorithm's designer.

[23] In appendix G we present two pieces of evidence to support this claim. First, both increases and decreases in the recommendation induce managers to update, and they do so in a similar manner. Second, the elasticities of manual price updates to both increases and decreases in the recommendation are similar. Moreover, we also find that, conditional on a price update, the manager copies the recommendation with a probability of 0.941 following an increase in the recommendation and with a probability of 0.944 following a decrease in the recommendation.

algorithm.[24] Finally, we discuss in appendix H other alternative models and demonstrate that they fail to account for the empirical patterns we observe in the data.

5.2 Analysis

We will next describe a set of theoretical results from the model which will then be matched with the stylized facts from the data. All proofs are relegated to appendix B. We begin our analysis with the problem of the hotel manager. Upon observing (r, c), she decides whether to initiate the information-acquisition process or maintain the current price. In the latter case, she expects a loss of $(\tilde{x}(r))^2 + \sigma_y^2 + \sigma_z^2$, where $\tilde{x}(r)$ is her belief about x given the observed recommendation r. In case she decides to continue her information acquisition, she expects a loss $l(r, c)$.

To characterize the expected loss $l(r, c)$, we need to calculate the hotel manager's payoff once she acquires information y. In this case, she will choose to copy the recommendation as along as $(r - \tilde{x}(r) - y)^2 + \sigma_z^2 < c_2$. Let $Y_0(r, c)$ denote the set of values of y for which this inequality holds. Notice also that $Y_0(r, c)$ is an interval centered at $r - \tilde{x}(r)$. The expected loss is then,

$$l(r, c) = c_1 + \int_{y \in Y_0(r,c)} \left((r - \tilde{x}(r) - y)^2 + \sigma_z^2\right) d\Psi_y(y) + \int_{y \notin Y_0(r,c)} c_2 d\Psi_y(y),$$

where Ψ_y is the cumulative distribution function of a mean-zero normal distribution with variance σ_y^2. The first lemma shows that larger costs and changes in the recommendation increase the hotel manager's expected loss when continuing the information acquisition past the status quo.

Lemma 1. *The continuation loss function $l(r, c)$ satisfies $l(r, c) = l(-r, c)$, increasing in c and non-decreasing in $|r|$. Furthermore,*

$$0 \leq l_r(r, c) \leq 2r(1 - \tilde{\lambda})^2 \int_{y \in Y_0(r,c)} d\Psi_y(y) \text{ for all } r > 0.$$

If $\sigma_y > 0$ and $\tilde{\lambda} < 1$, the loss function is strictly increasing in $|r|$ and both inequalities are strict.

[24] There is a rapidly growing literature on platform pricing when sellers use pricing algorithms (Huang, 2022; Johnson et al., forthcoming) and on collusion using AI pricing algorithms (Calvano et al., 2020; Asker et al., 2023). The collusion literature includes studies on how outsourcing pricing algorithms to a third party (Harrington Jr, 2022), or increased algorithmic forecasting accuracy (Miklós-Thal and Tucker, 2019), affects competition. Leisten (2022) investigates the effect of algorithmic pricing on competition when there is human override.

Because $l(r, c)$ is continuous and increasing in c, for every r there exists cost realization $c(r)$ such that a hotel manager continues to acquire information for costs lower than $c(r)$ and prefers to keep the current price for costs larger than $c(r)$. Notice that $c(r)$ is an even function. Conversely, for a given cost, the set of recommendations such that an update occurs can be written as the union of two intervals $(-\infty, -\bar{r}(c))$ and $(\bar{r}(c), \infty)$, for some function $\bar{r}(c)$. The following lemma provides a condition such that that both $\bar{r}(c)$ and $c(r)$ are strictly increasing functions.

Lemma 2. *Suppose that* $\frac{1}{2} \leq \tilde{\lambda} < 1$ *and* $\sigma_y > 0$*. Then* $\bar{r}(c)$ *is strictly increasing.*

The lemma says that higher cost realizations require larger deviations in the recommendation before the hotel manager considers re-evaluating the current price. Notice then that the probability of choosing a price that exactly matches the recommendation depends on the combination of two forces. First, the mass at $Y_0(r, c)$ is decreasing in $|r|$. This shows that, conditional on an update of the price, the probability of departing from the recommendation is higher at lower values of $|r|$, contradicting the empirical observations described above. Importantly, however, the probability of a price update is increasing in $|r|$. This implies that larger changes in the recommendation are associated with a higher chance that the hotel manager considers copying the recommendation in the first place. The resulting relationship between the size of the recommendation change and the probability of copying depends on the relative strength of these two forces. Let $\mu(r)$ denote the likelihood of matching the recommendation conditional on an update. The following proposition characterizes these effects formally.

Proposition 1. *If* $b_2\sigma_y^2 < b_1\sigma_z^2$*, then* $\mu(r) = 0$ *for all* r*. Else,* $\mu(0) > 0$*,* $\lim_{r\to 0} \mu(r) = 0$*, and there exists some* $r^* > 0$ *such that* $\mu_r(r) \geq 0$ *for all* $r \in (0, r^*)$*. In addition,* $\mu(r)$ *is (weakly) increasing if* $\tilde{\lambda} = 1$ *and (weakly) decreasing if* $F(c(0)) = 1$*.*

The two special cases highlighted in the proposition are instructive. First, if $\lambda = 1$, higher values of r induce hotel managers with higher opportunity costs of time to put attention to the price. This translates into a higher likelihood of copying the recommendation since there is no additional selection. In other words, a model with unbiased advice the hotel manager will be more likely to copy the recommendation, the bigger is the change in the recommendation. Second, if the hotel manager updates prices in every period (i.e. if $F(c(0)) = 1$), there is no inertia and higher r makes copying the recommendation less appealing because it is associated with a higher loss. This type of comparative static holds in models of strategic delegation where the hotel manager, the principal, is more likely to rubber-stamp *low* recommendations from the agent (Aghion and Tirole, 1997).

We now focus on the distribution of prices conditional on a departure from the recommendation. The expectation of such a price can be written as

$$E(p \mid r, y \notin Y_0(r,c)) = \tilde{x}(r) + E(y \mid r, y \notin Y_0(r,c)). \tag{1}$$

Since, $Y_0(r,c)$ is centered at $r - \tilde{x}(r)$, $E(y \mid r, y \notin Y_0(r,c))$ depends on $\tilde{\lambda}$. If $\tilde{\lambda} = 1$, then $r - \tilde{x} = 0$, $Y_0(r,c)$ is centered at the origin and hence $E(p \mid r, y \notin Y_0(r,c)) = \tilde{x}(r) = r$. Instead, if $\tilde{\lambda} \in (0.5, 1)$, then the conditional covariance of $(\tilde{x}, y)$ given that $y \notin Y_0(r,c)$ is negative, resulting in a dampening of the pass-through rate below $\tilde{\lambda}$.

Proposition 2. *The expected price conditional on the price departing from the recommendation satisfies*

$$E(p \mid r, y \notin Y_0(r,c)) \leq \tilde{\lambda}r, \ \textit{for all } r > 0$$

and

$$E(p \mid r, y \notin Y_0(r,c)) \geq \tilde{\lambda}r, \ \textit{for all } r < 0$$

with strict inequalities whenever $\sigma_y^2 > 0$ *and* $0.5 < \tilde{\lambda} < 1$*.*

The proposition implies that there is negative selection in unobservables and we cannot directly identify $\tilde{\lambda}$ from the pass-through rate.

Corollary 1. *Conditional on the hotel manager not copying the recommendation but updating the price, her private information is negatively correlated with the recommendation, i.e.*

$$Cov(y, r \mid y \notin Y_0(r,c)) \leq 0.$$

We finally address the problem of the algorithm. The algorithm chooses λ to maximize expected profits but this 'bias factor' is not directly observed by the hotel manager. An equilibrium is a triple $(\lambda, c(r), Y_0(c,r))$ such that λ maximizes profits given $(c(r), Y_0(c,r))$ and $(c(r), Y_0(c,r))$ are optimal given $\tilde{x} = \lambda r$. In general, a marginal increase in λ brings about three changes in the distribution of prices. First, it leads to a reduction in the variance in the distribution of recommendations, which necessarily induces the hotel manager to change prices less frequently. Second, it has an ambiguous impact on the probability that the hotel manager chooses a price that exactly matches the recommendation, because the function $\mu(r)$ is non-monotone. Third, it reduces the distance between the recommendation and the optimal price which translates directly into increased profits.

6 Estimation and Results

For the empirical implementation of the price-setting model in section 5, we assume that the information-acquisition cost c of the hotel manager follows a lognormal distribution with parameters $(0, \sigma_c)$. We have 6 structural parameters of which three govern the informational environment $(\sigma_x, \sigma_y, \sigma_z)$ and the other three capture the distribution of shocks (σ_c, b_1, b_2). Additionally, we need to infer the reduced-form parameter $\tilde{\lambda}$ which measures the equilibrium bias in the algorithmic recommendation.[25] To estimate these parameters, we use a method of simulated moments (SMM), minimum distance estimator with seven target moments that additionally imposes the restriction that there is no (secret) profitable deviation from the recommendation for the algorithm.[26] Four of these target moments depend directly on the joint distribution of recommendation and price updates, see rows 1 to 4 in Table 5. In addition, we use the likelihood of the updated price matching the recommendation exactly, both unconditionally and conditionally on the recommendation change exceeding 10% as well as the average price update rate, see rows 5 to 7 in Table 5.

Our estimation algorithm is implemented as follows. We first fix a level of the recommendation bias $\tilde{\lambda}$ and simulate the model to find structural parameter values that minimize the distance between the simulated moments and their observed targets. We then check whether a local deviation from $\tilde{\lambda}$ increases the revenue management company's payoff. If such a profitable deviation $\tilde{\lambda}'$ exists, we pick it as the new starting value and re-estimate the structural parameters. We repeat this process until we find a $\tilde{\lambda}$ and a set of distance minimizing parameters such that no profitable local deviations exist. We also try multiple starting values. In the case the algorithm finds two different equilibria with different parameter configurations, we choose the one with the smallest distance between simulated moments and target moments. The model is estimated both for the pooled data and for each hotel individually.[27]

To discuss identification, it is instructive to consider a special case of the model in which $\sigma_y = 0$.[28] Because there is no selection into updating based on payoff-relevant information,

[25] To identify $\tilde{\lambda}$ for a hotel, we assume that a given hotel always plays the same equilibrium with the recommendation algorithm, but allow for potentially different equilibria across hotels.

[26] The objective function depends on second moments and may not be quasiconcave. To search for a global optimum, we initiate the search procedure at multiple starting values, corresponding to the initial estimates we obtained for each hotel from a common initial guess. As a further robustness, we compare the results from the baseline estimation to that of the SMM estimation (with the same starting values) where we treat λ as an additional parameter to be estimated but do not impose any equilibrium conditions: they both yield similar values for the bias; see Figure 6, and Tables 7 and 12.

[27] A more detailed description of the estimation algorithm can be found in appendix I.

[28] A proof of identification for this case, as well as a discussion of why this particular assumption is unlikely to hold, is provided in appendix D.

Table 5: Targets for Pooled Data

Moment	Data	Model
$\sqrt{Var(p\|\text{Update})}$	0.074	0.074
$\sqrt{Var(r\|\text{Update})}$	0.068	0.067
$\sqrt{Var(p-r\|\text{Update})}$	0.035	0.035
$\sqrt{E(p \cdot r\|\text{Update})}$	0.068	0.066
$\Pr(\text{Copy Rec}\|\text{Update})$	0.840	0.842
$\Pr(\text{Copy Rec}\|\text{Update, Large Rec})$	0.947	0.951
$\Pr(\text{Update Rate})$	0.038	0.040

Notes: The first two rows report the standard deviation of the price (p) and the recommendation (r), both conditional on an Update. The third row reports the standard deviation of the difference between the price and the recommendation and the fourth reports the square root of the covariance (both variables have zero mean), all conditional on an update. Rows five and six report the rate of copying recommendations, both unconditionally and conditional on the recommendation change exceeding 10% (Large Rec). The last row reports the unconditional update rate.

the difference between the price and the recommendation directly determines the standard deviation of z, the covariance between r and p directly pins down λ, and the standard deviation of r determines the standard deviation of x for a given λ. Likewise, the ratio of the copy rate for large changes in the recommendation over the average copy rate determines the standard deviation of the cost distribution. The two remaining parameters can be directly obtained by matching the update rate and the average copy rate. While things are more complicated when $\sigma_y \neq 0$, each of these parameters is closely linked to the corresponding moment, with the standard deviation of price changes now helping to determine the non-zero σ_y.

The first row of Table 6 presents the results of the estimation routine on the pooled dataset, along with the bootstrapped standard errors. We find that the private information of the algorithm accounts for less than 20% of the total variance of the optimal price. This means that hotel managers' private information is at least five times as valuable as that of

Table 6: Parameter Estimates of Model

Hotel	σ_x	σ_y	σ_z	σ_c	b_1	b_2	λ
Pooled	0.038	0.053	0.019	2.28	0.090	0.897	0.830
	(0.001)	(0.001)	(0.003)	(0.03)	(0.015)	(0.029)	
A	0.032	0.092	0.126	1.35	0.021	0.802	0.865
	(0.002)	(0.005)	(0.011)	(0.03)	(0.012)	(0.024)	
B	0.023	0.056	0.024	2.06	0.032	0.644	0.840
	(0.001)	(0.002)	(0.008)	(0.05)	(0.011)	(0.028)	
C	0.026	0.034	0.026	2.25	0.039	0.383	0.835
	(0.001)	(0.001)	(0.001)	(0.02)	(0.006)	(0.017)	
D	0.017	0.037	0.014	2.64	0.035	0.676	0.690
	(0.001)	(0.001)	(0.003)	(0.06)	(0.012)	(0.027)	
E	0.028	0.032	0.040	1.79	0.016	0.273	0.740
	(0.001)	(0.001)	(0.002)	(0.02)	(0.002)	(0.009)	
F	0.019	0.036	0.001	1.84	0.027	0.397	0.660
	(0.001)	(0.001)	(0.005)	(0.04)	(0.007)	(0.025)	
G	0.032	0.026	0.035	2.54	0.078	0.809	0.715
	(0.001)	(0.001)	(0.002)	(0.03)	(0.008)	(0.025)	
H	0.036	0.064	0.092	1.43	0.010	0.681	0.600
	(0.002)	(0.006)	(0.007)	(0.03)	(0.005)	(0.032)	
I	0.034	0.083	0.08	1.54	0.034	0.644	0.930
	(0.002)	(0.003)	(0.005)	(0.03)	(0.011)	(0.024)	

Notes: Estimated parameter values for each hotel, A to I, and pooled across hotels. Given our parametrization, the mean of the distribution is $\exp(\ln b_1 + \sigma_c^2/2)$. Bootstrapped standard errors in parenthesis. The reported λ is the computed loss-minimizing recommendation bias satisfying the equilibrium condition.

the algorithm. However, accessing this information requires substantial effort by the manager. Figure 15 in the appendix shows the distribution of the costs associated with copying recommendations and manually setting prices. These cost differences imply a considerable dispersion between actual prices and counterfactual optimal prices. This disparity also implies that adjustment costs do rather reflect costly managerial attention than fear of consumer backlash (Rotemberg, 2005).

Our results suggest a significant bias in recommendations with $\lambda = 0.83$ when pooling over all hotels. Because most price updates match the recommendation exactly, this bias implies biased (sub-optimal) prices. Nevertheless, the welfare impact of the recommenda-

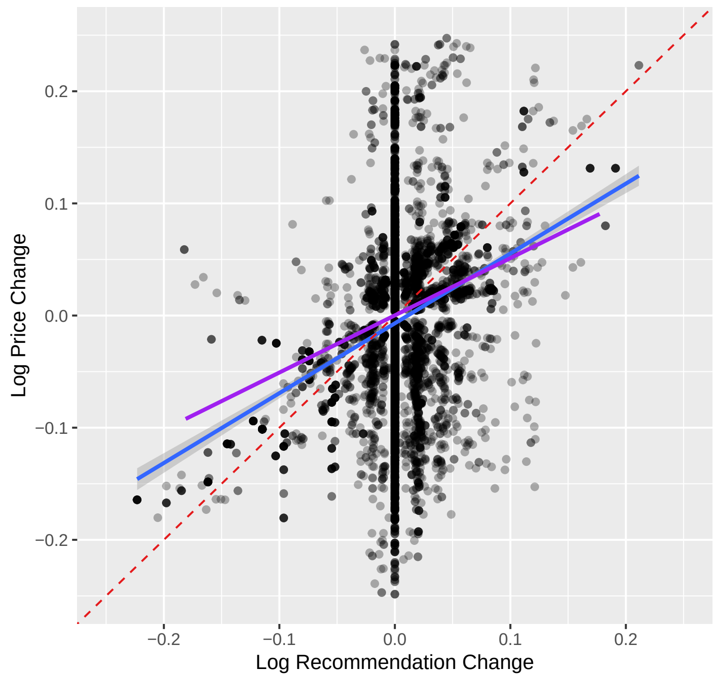

Figure 5: Model Fit: Recommendations and Prices

Notes: Each point represents a price update that does not match the recommendation. The horizontal axis is the log change in the recommendation and vertical axis is the log change in the price. The blue line shows a linear fit of the data with a 95% confidence interval, the purple line shows a linear fit to simulated data with our estimated parameters and the dashed red line plots the 45-degree line.

tion bias is ameliorated by the fact that the manager selects into the decision to match the recommendation, see Proposition 2.

Table 5 also shows that the model is able to fit the target moments well. It also does a reasonably good job at replicating the empirical facts regarding the relationship between price updates and recommendations described in section 4. For instance, it predicts an update rate of about 15% when the recommendation exceeds 5%, which is slightly higher, but reasonably close to the data. It also generates a relation between recommendations and prices, conditional on observing a price update, that it is consistent with the data, see Figure 5.

We further evaluate the fit of the model by performing an alternative estimation procedure in which we take λ as a primitive of the data and minimize the same distance estimator with seven targets and seven moments. The estimated parameters are similar and, in particular, the estimated λ across different hotels, while somewhat larger, is highly correlated with the one obtained from the baseline estimation as shown in Figure 6.

We then run our estimation routine separately for each hotel and report the hotel-level

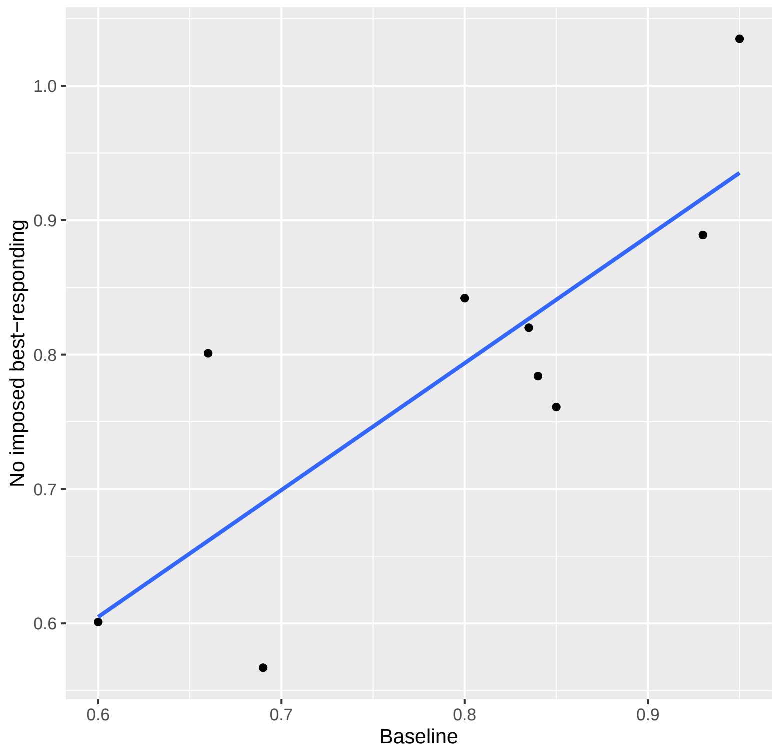

Figure 6: Bias Parameter Validation

Notes: Each point represents a hotel. The horizontal axis gives the bias parameter identified using the optimality condition for the algorithm and the vertical axis gives the bias parameter using only the hotel manager's information.

results. Most hotels are accurately represented by the pooled data as shown in Table 6. The variance of the recommendation's information, σ_x, accounts for around 20-30% of the total variance for all hotels. There is, however, considerable heterogeneity in the precision of the freely available information (measured by σ_y) relative to the total information that is potentially available to the manager ($\sigma_y + \sigma_z$). The strategic bias in the recommendations varies across hotels between 0.60 (hotel H) and 0.93 (hotel I). The differences are driven by the information held by the algorithm and by heterogeneity in the level of managers' inertia, with higher inertia implying larger gains from delegation.

7 Counterfactuals

We are now in a position to investigate whether hotels would gain from delegating pricing decision to the algorithm. Delegation to a fully automated algorithm offers several advantages.

Table 7: Counterfactuals

Hotel	Benchmark	Profit Loss	Delegation	Biased	No Rec
A	0.998	0.991	0.961	0.961	0.999
B	0.990	0.977	0.875	0.879	0.998
C	0.984	0.967	0.727	0.734	0.998
D	0.984	0.970	0.835	0.851	0.996
E	0.986	0.970	0.761	0.779	0.999
F	0.996	0.990	0.774	0.800	1.000
G	0.984	0.968	0.643	0.673	0.999
H	0.998	0.991	0.909	0.924	0.999
I	0.994	0.983	0.920	0.921	0.998

Notes: The value in the first column (Benchmark) corresponds to the welfare loss in the status quo relative to the welfare loss under complete inaction. The value in the second column (Profit Loss) is the implied accounting profit loss, disregarding adjustment costs, relative to inaction. The third column (Delegation) represents the welfare loss in the counterfactual exercise of full delegation to the algorithm, again relative to inaction. The fourth column (Biased) describes the expected welfare loss from a counterfactual where the decision is delegated to the algorithm which continues to produce biased recommendations relative to complete inaction. The last column (No Rec) shows the expected welfare loss if the hotel manager has no access to the recommendation or copying it, again relative to inaction.

It eliminates adjustment costs for hotel managers and thereby also delays in decision-making. Delegation also leads to truthful recommendations by the algorithm as it eliminates the strategic conflict between the algorithm and the manager caused by the latter's inertia. A potential downside of full delegation is that the algorithm does not have access to the manager's informative signals y and z, which are even more valuable than the algorithm's signal x according to our data. The main insight, however, is that the managerial inertia reveals that these informative signals come with a significant cost for the manager, greatly reducing their value for human decision making.

To conceptualize our counterfactual experiments, we first compute the expected loss in profit for each hotel under the status quo setting, in which the manager sets prices given the algorithmic recommendations, relative to the profit loss they would experience if they would never update their prices (complete inaction). This metric is independent of parameter η and takes into consideration that some hotels experience a more volatile environment than others. Formally, this loss in profit is given by

$$w_i = \frac{1}{\sigma_x^2 + \sigma_y^2 + \sigma_z^2} \int \left(\int_0^{c(\frac{x}{\lambda})} l\left(\frac{x}{\lambda}, c\right) dF(c) + \int_{c(\frac{x}{\lambda})}^{\infty} (x^2 + \sigma_y^2 + \sigma_z^2) \right) d\Psi(x). \tag{2}$$

The denominator here is simply the maximum utility loss from inaction. The first term in the numerator is the expected loss (including adjustment costs) when the hotel manager either copies the recommendation or adjusts manually.[29] The second term is the expected loss in case of optimal inaction, given the realization of the recommendation and the adjustment costs. Notice that lower values correspond to more efficient outcomes, with $w_i = 0$ being the first-best outcome. The estimate of equation (2) for each hotel in our sample is reported in the Benchmark column in Table 7.

We can also directly compute the residual losses under delegation based on the parameter estimates in section 6. We consider two extreme scenarios. First, assume that the algorithm lacks any incentives to distort its recommendations and consequently its pricing decision under full delegation (Delegation). We then have $p = x$, which yields a loss in profit of

$$w_i = \frac{\sigma_y^2 + \sigma_z^2}{\sigma_x^2 + \sigma_y^2 + \sigma_z^2}.$$

Alternatively, assume that the algorithm does not fully re-optimize its recommendations but instead continues to misrepresent its information (Biased). In this scenario, we have $p = x/\lambda$ and a profit loss of

$$w_i = \frac{(1-\lambda)^2}{\lambda^2} \frac{\sigma_x^2}{\sigma_x^2 + \sigma_y^2 + \sigma_z^2}.$$

Our counterfactual estimation results for the different scenarios are summarized in Table 7. The results illustrate that the status quo (Benchmark column in Table 7) is only slightly better than complete inaction while delegating pricing to the algorithm (Delegation column) would do much better even if the algorithm's prices were biased as in the status quo (Biased column). More specifically, we find that the status quo can mitigate only $1-3\%$ of the profit loss from complete inaction across all hotels (Profit Loss). These gains are even smaller if the adjustment and information acquisition costs of managers to achieve these gains are taken into account (Benchmark). These findings align with the empirical observation that managers

[29] The function $l(\cdot)$ is formally defined in section 5.2.

infrequently adjust prices. And when they do, they update prices with delay and usually just copy the recommendation.

Delegating pricing to the algorithm is likely to improve outcomes considerably, see Table 7 (Delegation column). Results show that a hotel that fully delegates pricing to an unbiased algorithm would see a reduction of 4 to 36 percent in losses accrued from mispricing relative to inaction. From column Biased in the table it can be seen that the welfare loss for each hotel is only slightly larger when the algorithm uses biased recommendations instead of the unbiased ones under full delegation.[30] Taken together, about 80% of the gains from delegation come from the algorithm adjusting prices much more frequently, 10% depends on the algorithm reporting truthfully, and the remaining 10% result from costless information processing. Moreover, the potential improvements from delegation to the algorithm varies strongly across hotels as shown in Table 7. For hotels A, H and I, for example, delegating pricing to the algorithm would leave significant surplus on the table because our estimates suggest that most of the variation in optimal pricing can only be discovered by the local hotel manager.[31]

Lastly, the profit loss hotel managers would accrue if they were not able to access the algorithm's recommendation at all is reported in column No Rec in Table 7. For all hotels, the resulting payoff losses are nearly identical to the baseline of complete inaction, again highlighting how infrequent the manual updates are relative to copying recommendations and periods of inaction. In other words, the adjustment costs are typically so high that even when the hotel manager manually updates the price, updating costs eat up a lion's share of the accrued benefit. The algorithm hence provides clear value by offering a cheap way of adjusting prices more frequently and by incorporating additional price information.

As AI technologies and data availability advance, it is likely that the recommendation algorithm's share of the totally available price information will continue to grow. Alternatively, hotel managers could respond by investing in their own market analysis to increase their share of the total information. It is therefore natural to ask how a hotel manager's payoff varies with her share of the totally available price information. Figure 7 plots the hotel manager's payoff in the pooled sample from the current interaction with the recommendation algorithm as well as from delegation to the algorithm, both as a function of the share of information

30 We report appendix I in the results for an alternative estimation method. Overall, results are similar with gains ranging from 5 to 40%.

31 Note that the above results are likely a lower bound on the total increase in profits from delegation, as it may induce "coordinated effects" on rival firms who benefit from the increased sensitivity of prices to market conditions (see, Harrington Jr, 2022).

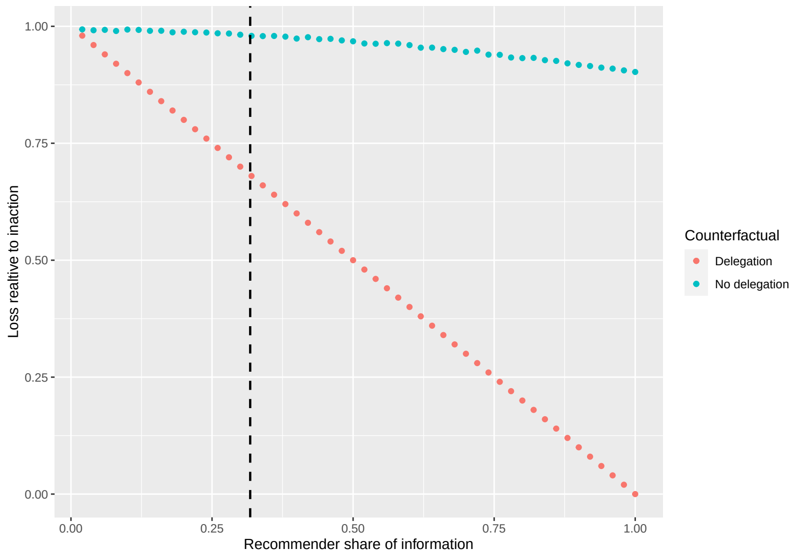

Figure 7: Counterfactual Information Shares and Hotel Manager Payoffs

Notes: The figure shows how the hotel manager's payoff changes when the recommendation algorithm's share of the total information changes for the pooled sample. We plot the payoff losses relative to complete inaction for both the current environment and a counterfactual where pricing is delegated to the recommendation algorithm. The dashed vertical line shows the status quo. The model parameters are the ones estimated in the first row of Table 6. We re-estimate a new equilibrium bias for each new division of information, see Figure 9 in appendix F.

held by the recommendation algorithm.

We re-estimate the hotel manager's payoff in these two scenarios for different values of $\sigma_x(\alpha)^2 = \alpha(\sigma_x^2 + \sigma_y^2 + \sigma_z^2)$ by varying the algorithm's share of the information, $\alpha \in (0, 1)$, and holding the variance of optimal prices, $\sigma_x(\alpha)^2 + \sigma_y(\alpha)^2 + \sigma_z(\alpha)^2$, and the relative importance of y and z, $\frac{\sigma_y(\alpha)^2}{\sigma_z(\alpha)^2}$, constant at the levels estimated on the first row of Table 6. We also re-estimate the equilibrium bias of the recommendation algorithm, for each possible share of the totally available price information, and report it in Figure 9 in appendix F. The remaining parameters are taken as in the first row of Table 6. The dashed vertical line in Figure 7 plots the share of information held by the recommendation algorithm we estimated for the status quo setup ($\alpha = \frac{\sigma_x^2}{\sigma_x^2+\sigma_y^2+\sigma_z^2}$).

We find that when the algorithm's relative share of information increases, the hotel manager becomes better off in the current 'status quo' setting because expensive manual adjustments are less necessary and copying recommendations leads to more accurate and profitable pricing. Even if the algorithm would get much better, the hotel manager's adjustment costs still prevent her from realizing a large part of the potential gains under the status quo (see No Delegation in Figure 7). In addition to the adjustment cost of the manager, we also consider how the algorithm's biased recommendations influence the expected loss under the status quo in equilibrium. Figure 9 in appendix F illustrates that the payoff increases resulting from the larger information share of the algorithm are initially tempered by increases in the algorithm's equilibrium bias. When the algorithm holds next to no information, exaggerating its signals has a negligible impact on the hotel manager's update frequency. The impact of the biased exaggeration, however, increases in the algorithm's information and seems to plateau, or even decrease, when the algorithm holds most of the available information. This pattern is intuitive, as now, hotel managers copy recommendations most of the time and hence the cost of biasing those recommendations becomes higher. What is also striking about Figure 9 is that the current setup outperforms delegation only if the recommendation algorithm holds almost no pricing information. Again, a result of the relatively rare manual updates which implies that adjustment costs are very often prohibitively high.

Finally, our model also permits parameter configurations in which delegation is clearly inferior compared to the current setup. Figure 10 in appendix F illustrates the effects of an increase of σ_c, the standard deviation of the normal distribution that underlies the log-normal determining adjustment costs, on the benefits of delegation. The figure clearly shows that, if the hotel manager possesses enough information and her adjustment costs are low enough to permit for manual adjustments, delegation is clearly inferior to the current setting in which algorithmic recommendations are used to augment hotel managers' pricing decisions.

8 Conclusion

Algorithmic recommendations are used extensively to support decision making in organizations. In this paper, we provide a framework for understanding the strategic interaction between algorithmic recommendations and human decisions. The crucial friction in our model originates in managerial inattention, leading to biased communication and pricing decisions. Applying our model of information processing to a dataset containing millions of hotel-room price recommendations, we demonstrate that full delegation to the algorithm is likely to be welfare-improving, even if it forgoes the potential benefits of richer information.

Our findings point to a novel element in the intricate relationship between algorithmic advisors and human decision makers. Previous work has studied the impact of heterogeneous preferences and skills, as well as potential bias in the processing of algorithmic advice by humans. We show that humans may become a bottleneck in the decision making process as they struggle to keep up with the arrival of frequently changing information, thereby severely limiting the benefits of advice. In particular, we argue that when a human decision maker has access to an easy default or status quo option (e.g. 'keep current price') and deviating from that option (e.g. 'update price') incurs adjustment costs, this can lead to severe conflict of interest between the decision maker and the algorithm's designer who does not incur such costs. Examples of such settings include diagnosis recommendations for doctors, recommendations systems for parole decision (Berk, 2017), monitoring adherence to government regulation (Glaeser et al., 2021) and restocking inventory (Shang et al., 2008).[32][33]

In general, an algorithmic recommendation has three potential benefits: it provides additional information, it simplifies the decision maker's task by offering an easily selectable alternative and it may also redirect the attention of the decision maker to carefully think about other alternatives, leading to better decisions as a consequence. These changes typically further benefit the algorithm's designer and hence the slow pace at which the decision maker takes actions becomes even less optimal for the algorithm's designer and may therefore incentivize even more biased recommendations.

Ludwig and Mullainathan (2021) argue that even best-practice algorithmic design has been unable to efficiently incorporate both preferences and information of human decision makers into recommendation algorithms (also known as the *override problem*). Our work demonstrates that the problem can be further complicated by strategic considerations. In our setting, human actors who perceive recommendations as distorted, strategically counterbalance those distortions. Responding strategically to algorithmic recommendations can be especially important for human decision makers in judicial decisions (Kleinberg et al., 2018) or hiring decisions (Hoffman et al., 2018) where the designer of the algorithm would like to correct for underlying human biases, while the biased decision maker may have incentives to strategically 'correct' the recommendation given that she understands that the algorithmic recommendation is attempting to de-bias her decisions.

[32] A doctor, for instance, can easily default to the most common cause of prominent symptoms when making a diagnosis, government officials can choose to not check adherence to regulation, parole boards can keep an inmate locked in, and a store manager can decline to restock their inventory.

[33] Even though hotel pricing decisions may seem to have lower stakes relative to decisions made by judges and doctors, a growing literature, especially from courts, suggests that these high-stake decisions are affected by arguably much more irrelevant outside factors than effort costs. Examples include the fate of the presiding judge's hometown football team (Eren and Mocan, 2018), temperature outside the courtroom (Heyes and Saberian, 2019) or unrelated previous judicial decisions (Chen et al., 2016a).

There are many potential avenues for future research. An obvious extension of the present paper would involve an explicit, fully dynamic model of price adjustment. The challenge here is to handle strategic communication when managers may be tempted to wait for further information before acting. Another question that we have not attempted to answer is why recommendation systems are not substituted with delegation to the algorithmic advisor, even when this would potentially benefit both economic agents. The answer may have to do with the perception that algorithmic systems are biased, as in our case, and that they are likely to make costly mistakes. For example, Dietvorst et al. (2015) argues that human decision makers have a low tolerance for machine errors, and would rather rely on less precise human advice. Relatedly, in a recommendation system the responsibility for mistakes typically rests with the final decision maker while the designer of the algorithm largely escapes liability for poor advice. This is likely to be a significant reason for using recommendation systems, especially in revenue management and other economic consulting where disentangling the effect of poor pricing advice from poor general management can be difficult. Finally, it would be interesting to study strategic communication in environments in which the decision problem is better described as a prediction problem and the researcher has access to the ex-post optimal choice. This would allow to directly measure the degree of bias in communication, rather than relying on the equilibrium response of the decision maker, thereby enabling model validation.

References

Abada, Ibrahim and Xavier Lambin, "Artificial intelligence: Can seemingly collusive outcomes be avoided?," *Management Science*, 2023.

Aghion, Philippe and Jean Tirole, "Formal and real authority in organizations," *Journal of Political Economy*, 1997, *105* (1), 1–29.

Agrawal, Ajay, Joshua S Gans, and Avi Goldfarb, "Exploring the impact of artificial intelligence: Prediction versus judgment," *Information Economics and Policy*, 2019, *47*, 1–6.

Alvarez, Fernando E., Francesco Lippi, and Luigi Paciello, "Optimal price setting with observation and menu costs," *Quarterly Journal of Economics*, 2011, *126* (4), 1909–1960.

Asker, John, Chaim Fershtman, and Ariel Pakes, "The impact of artificial intelligence design on pricing," *Journal of Economics & Management Strategy*, 2023.

Assad, Stephanie, Robert Clark, Daniel Ershov, and Lei Xu, "Algorithmic pricing and competition: Empirical evidence from the German retail gasoline market," *Journal of Political Economy*, forthcoming.

Athey, Susan C., Kevin A. Bryan, and Joshua S. Gans, "The allocation of decision authority to human and artificial intelligence," *AEA Papers and Proceedings*, 2020, *110*, 80–84.

Backus, Matthew, Thomas Blake, and Steven Tadelis, "On the empirical content of cheap-talk signaling: An application to bargaining," *Journal of Political Economy*, 2019, *127* (4), 1599–1628.

Berk, Richard, "An impact assessment of machine learning risk forecasts on parole board decisions and recidivism," *Journal of Experimental Criminology*, 2017, *13* (2), 193–216.

Bloedel, Alexander W. and Ilya Segal, "Persuading a Rationally Inattentive Agent," 2020. working paper.

Bolton, Gary E and Elena Katok, "Learning by doing in the newsvendor problem: A laboratory investigation of the role of experience and feedback," *Manufacturing & Service Operations Management*, 2008, *10* (3), 519–538.

_, Axel Ockenfels, and Ulrich W Thonemann, "Managers and students as newsvendors," *Management Science*, 2012, *58* (12), 2225–2233.

Brown, Zach Y. and Alexander MacKay, "Competition in pricing algorithms," *American Economic Journal: Microeconomics*, 2023, *15* (2), 109–156.

Bundorf, M Kate, Maria Polyakova, Cheryl Stults, Amy Meehan, Roman Klimke, Ting Pun, Albert Solomon Chan, and Ming Tai-Seale, "Machine-based expert recommendations and insurance choices among Medicare Part D enrollees," *Health Affairs*, 2019, *38* (3), 482–490.

Calvano, Emilio, Giacomo Calzolari, Vincenzo Denicolo, and Sergio Pastorello, "Artificial intelligence, algorithmic pricing, and collusion," *American Economic Review*, 2020, *110* (10), 3267–97.

Camara, Fanny and Nicolas Dupuis, "Structural estimation of expert strategic bias: The case of movie reviewers," 2014. working paper.

Caro, Felipe and Anna Sáez de Tejada Cuenca, "Believing in analytics: Managers' adherence to price recommendations from a DSS," *Manufacturing & Service Operations Management*, 2023, *25* (2), 524–542.

Chan, David C., Matthew Gentzkow, and Chuan Yu, "Selection with variation in diagnostic skill: Evidence from radiologists," *Quarterly Journal of Economics*, 2022.

Chen, Daniel L, Tobias J Moskowitz, and Kelly Shue, "Decision making under the gambler's fallacy: Evidence from asylum judges, loan officers, and baseball umpires," *Quarterly Journal of Economics*, 2016, *131* (3), 1181–1242.

Chen, Le, Alan Mislove, and Christo Wilson, "An empirical analysis of algorithmic pricing on Amazon marketplace," in "WWW'16: Proceedings of the 25th International Conference on World Wide Web" 2016, pp. 1339–1349.

Cowgill, Bo and Megan T. Stevenson, "Algorithmic social engineering," *AEA Papers and Proceedings*, 2020, *110*, 96–100.

Crawford, Vincent P. and Joel Sobel, "Strategic information transmission," *Econometrica*, 1982, pp. 1431–1451.

Currie, Janet and W. Bentley MacLeod, "Diagnosing expertise: Human capital, decision making, and performance among physicians," *Journal of Labor Economics*, 2017, *35* (1), 1–43.

Dean, Mark, Özgür Kıbrıs, and Yusufcan Masatlioglu, "Limited attention and status quo bias," *Journal of Economic Theory*, 2017, *169*, 93–127.

DellaVigna, Stefano and Matthew Gentzkow, "Persuasion: Empirical evidence," *Annual Reviews of Economics*, 2010, *2* (1), 643–669.

Dietvorst, Berkeley J., Joseph P. Simmons, and Cade Massey, "Algorithm aversion: People erroneously avoid algorithms after seeing them err," *Journal of Experimental Psychology: General*, 2015, *144* (1), 114.

Eren, Ozkan and Naci Mocan, "Emotional judges and unlucky juveniles," *American Economic Journal: Applied Economics*, 2018, *10* (3), 171–205.

Garcia, Daniel, Juha Tolvanen, and Alexander K. Wagner, "Dynamic demand estimation using managerial responses to automated price recommendations," *Management Science*, 2022, *68* (11), 7918–7939.

Glaeser, Edward, Andrew Hillis, Hyunjin Kim, Scott Duke Kominers, and Michael Luca, "Decision authority and the returns to algorithms," 2021. working paper.

Hansen, Karsten T, Kanishka Misra, and Mallesh M Pai, "Frontiers: Algorithmic collusion: Supra-competitive prices via independent algorithms," *Marketing Science*, 2021, *40* (1), 1–12.

Harrington Jr, Joseph E, "The effect of outsourcing pricing algorithms on market competition," *Management Science*, 2022, *68* (9), 6889–6906.

Heyes, Anthony and Soodeh Saberian, "Temperature and decisions: evidence from 207,000 court cases," *American Economic Journal: Applied Economics*, 2019, *11* (2), 238–265.

Hoffman, Mitchell, Lisa B. Kahn, and Danielle Li, "Discretion in hiring," *Quarterly Journal of Economics*, 2018, *133* (2), 765–800.

Huang, Yufeng, "Pricing Frictions and Platform Remedies: The Case of Airbnb," 2022. Working paper.

Johnson, Justin P., Andrew Rhodes, and Matthijs R. Wildenbeest, "Platform design when sellers use pricing algorithms," *Econometrica*, forthcoming.

Kamenica, Emir, "Bayesian persuasion and information design," *Annual Review of Economics*, 2019, *11*, 249–272.

_ **and Matthew Gentzkow**, "Bayesian persuasion," *American Economic Review*, 2011, *101* (6), 2590–2615.

Kartik, Navin, Marco Ottaviani, and Francesco Squintani, "Credulity, lies, and costly talk," *Journal of Economic Theory*, 2007, *134* (1), 93–116.

Kleinberg, Jon, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan, "Human decisions and machine predictions," *Quarterly Journal of Economics*, 2018, *133* (1), 237–293.

Leisten, Matthew, "Algorithmic Competition, with Humans," 2022. working paper.

Ludwig, Jens and Sendhil Mullainathan, "Fragile Algorithms and Fallible Decision-Makers: Lessons from the Justice System," *Journal of Economic Perspectives*, 2021, *35* (4), 71–96.

Maćkowiak, Bartosz, Filip Matějka, and Mirko Wiederholt, "Rational inattention: A review," *Journal of Economic Literature*, 2023, *61* (1), 226–273.

Miklós-Thal, Jeanine and Catherine Tucker, "Collusion by algorithm: Does better demand prediction facilitate coordination between sellers?," *Management Science*, 2019, *65* (4), 1552–1561.

Nakamura, Emi and Jón Steinsson, "Five facts about prices: A reevaluation of menu cost models," *Quarterly Journal of Economics*, 2008, *123* (4), 1415–1464.

Ribers, Michael A. and Hannes Ullrich, "Machine learning and physician prescribing: a path to reduced antibiotic use," 2023. working paper.

Rotemberg, Julio J., "Customer anger at price increases, changes in the frequency of price adjustment and monetary policy," *Journal of Monetary Economics*, 2005, *52* (4), 829–852.

Schweitzer, Maurice E. and Gérard P. Cachon, "Decision bias in the newsvendor problem with a known demand distribution: Experimental evidence," *Management Science*, 2000, *46* (3), 404–420.

Shang, Jennifer, Pandu R Tadikamalla, Laurie J Kirsch, and Lawrence Brown, "A decision support system for managing inventory at GlaxoSmithKline," *Decision Support Systems*, 2008, *46* (1), 1–13.

Sobel, Joel, "Giving and receiving advice," in Daron Acemoglu, Manuel Arellano, and Eddie Dekel, eds., *Advances in Economics and Econometrics: Tenth World Congress*, Econometric Society Monographs, Cambridge University Press, 2013, chapter 10, pp. 305–341.

Supplementary Materials

for

"Strategic Responses to Algorithmic Recommendations: Evidence from Hotel Pricing"

by D. Garcia, J. Tolvanen, and A.K. Wagner

A Full Sample Pass-Through Rates

In addition to the pass-through regressions in Table 4, which are based on the main sample, we report for robustness the corresponding results for the full sample in Table 8. For the definition of the main sample and the full sample, as well as their summary statistics, see section 3 and Table 1 for details. The pass-through regressions for the full sample in Table 8 show an attenuation bias in actual price changes because of the measurement error introduced by recommendation-price histories which we exclude in the main sample. This implies that the strategic bias in recommendations estimated from the main sample would be even larger using the full sample.

Table 8: Pass-Through Rates of Recommendation (Full Sample)

	Change in actual price			
	All	Manually Updated		
Rec Change	0.781***	0.539***	0.535***	0.420***
	(0.005)	(0.013)	(0.013)	(0.028)
Days ahead Polynomial	No	No	Yes	Yes
Room × Month FE	No	No	No	Yes
N	130,669	130,669	130,669	130,669

Notes: Linear regression model using the full sample. The dependent variable is the cumulative change in the actual rate since the last price update. Recommendation is the cumulative (log) change in the recommendation since the last price update. All regressions include all updates. Coefficients for *Manually Updated* correspond to the interaction term of recommendation × manual. *Room* is the room type and *Month* refers to the arrival month. Significance levels: *** $p < 0.001$

B Proofs

Proof. Proof of Lemma 1. We first establish that $0 \leq l_r(r,c) \leq 2(1-\tilde{\lambda})^2 \int_{y\in Y_0(r,c)} d\Psi_y(y)r$, for $r > 0$ (and vice versa). Since the integrand of the first element is exactly equal to c_2 at the boundaries, only the derivative of the integrand matters. Hence,

$$l_r(r,c) = 2(1-\tilde{\lambda}) \int_{y\in Y_0(r,c)} (r - \tilde{x}(r) - y) d\Psi_y(y). \tag{3}$$

Consider any $r > 0$. Since $\tilde{x} = \tilde{\lambda} r$ and $Y_0(r,c)$ is an interval centered at $r - \tilde{x}(r)$, then the symmetry of the normal distribution about zero implies that $0 \leq \int_{y\in Y_0(r,c)} y d\Psi_y(y) \leq (1-\tilde{\lambda})r$. The inequalities are strict if $\sigma_y > 0$ and $\tilde{\lambda} < 1$. Substituting the end points of this interval into (3) yields both, that l is increasing in r for $r > 0$, and the first claimed inequality in the lemma. When $r < 0$, an analogous argument shows that $(1-\tilde{\lambda})r \leq \int_{y\in Y_0(r,c)} y d\Psi_y(y) \leq 0$ and hence the inequalities are reversed, which proves the second inequality in the lemma and that l is increasing in $|r|$.

Taking a derivative with respect to c we have simply $l_c(r,c) = b_2 \int_{y\notin Y_0(r,c)} d\Psi_y(y) > 0$. □

Proof. Proof of Lemma 2. The set of values of r and c for which the hotel manager is indifferent between keeping the old price and gathering gathering additional information is implicitly defined by the identity

$$(\tilde{\lambda} r)^2 + \sigma_y^2 + \sigma_z^2 \equiv l(r,c)$$

An application of the implicit function theorem to the positive root of this identity then implies that

$$\bar{r}_c(c) = \frac{l_2(\bar{r},c)}{2\tilde{\lambda}^2 \bar{r} - l_1(\bar{r},c)},$$

This is positive, since the numerator is positive and, when $\bar{r} > 0$, the denominator satisfies

$$\begin{aligned} 2\tilde{\lambda}^2\bar{r} - l_r(\bar{r},c) \quad &> \quad 2\tilde{\lambda}^2\bar{r} - 2\bar{r}(1-\tilde{\lambda})^2 \int_{y\in Y_0(\bar{r},c)} d\Psi_y(y) \\ &\geq \quad 2\tilde{\lambda}^2\bar{r} - 2(1-\tilde{\lambda})^2\bar{r} \geq 0, \end{aligned}$$

where the first inequality follows from the previous lemma, the last inequality from the assumption of $\tilde{\lambda} \geq \frac{1}{2}$. □

Proof. Proof of Proposition 1. Notice first that

$$\begin{aligned} &(r - \tilde{x}(r) - y)^2 + \sigma_z^2 < b_2 c \\ \Leftrightarrow \quad & x - r - \sqrt{b_2 c - \sigma_z^2} < y < x - r + \sqrt{b_2 c - \sigma_z^2}. \end{aligned}$$

Denote $d(c) := \sqrt{\max\{b_2 c - \sigma_z^2, 0\}}$. Then we can write

$$\mu(r) = \frac{\int_0^{c(r)} \left(\Psi_y(r - \tilde{x}(r) + d(c)) - \Psi_y(r - \tilde{x} - d(c))\right) dF(c)}{F(c(r))}.$$

If $b_2\sigma_y^2 < b_1\sigma_z^2$, then $b_2 c(0) < \sigma_z^2$, $d(c(0)) = 0$, and, therefore, the hotel manager will acquire signal z even if $r = 0$ and $y = 0$. Hence, $\mu(r) = 0$ for all r. Instead if $b_2\sigma_y^2 > b_1\sigma_z^2$, $d(c(0)) > 0$, and hence $\mu(0) > 0$. In addition, as $r \to \infty$, the integrand vanishes, while the denominator converges to 1. Hence, $\lim_{r\to\infty} \mu(r) = 0$. Finally, to see that $\mu(r)$ is increasing in a neighborhood of $r = 0$ observe that

$$\begin{aligned}
\mu_r(r) &= \frac{1}{F(c(r))}\left(F'(c(r))c'(r)\eta(r - \tilde{x}, d(c(r))) + 2\int_0^{c(r)} \eta_1(r - \tilde{x}, d(c))(1 - \tilde{\lambda})dF(c)\right) \\
&- \frac{1}{F(c(r))^2} F'(c(r))c'(r) \int_0^{c(r)} \eta(r - \tilde{x}, d(c(r)))dF(c) \\
&= \frac{F'(c(r))c'(r)}{F(c(r))^2} \int_0^{c(r)} \Big(\eta(r - \tilde{x}, d(c(r))) - \eta(r - \tilde{x}, d(c))\Big)dF(c) \\
&+ \frac{2(1 - \tilde{\lambda})}{F(c(r))} \int_0^{c(r)} \eta_1(r - \tilde{x}, d(c))dF(c),
\end{aligned}$$

with $\eta(a, b) = \Psi_y(a+b) - \Psi_y(a-b)$ is decreasing in a and increasing in b. Hence, the first term in the last step are weakly positive and the second is weakly negative. Since $c'(0) = 0$, both terms are zero at $r = 0$ and the sign of $\mu(r)$ depends on the second derivative. Disregarding terms that vanish at $r = 0$, we have

$$\begin{aligned}
\mu_{rr}(0) &= \frac{F'(c(0))c''(0)}{F(c(0))^2} \int_0^{c(0)} \Big(\eta(0, d(c(0))) - \eta(0, d(c))\Big)dF(c) \\
&= \frac{F'(c(0))c''(0)}{F(c(0))^2} \int_{\frac{\sigma_z^2}{b_2}}^{c(0)} \Big(\eta(0, d(c(0))) - \eta(0, d(c))\Big)dF(c) > 0,
\end{aligned}$$

by the assumption above. For $\lambda = 1$, the second term is zero and the first term is weakly positive so $\mu(r)$ is weakly increasing. If $F(c(0)) = 1$, the first term is always zero and hence $\mu(r)$ is weakly decreasing. □

Proof. Proof of Proposition 2. Assume first that $r > 0$. By (1), it is sufficient to establish that $\tilde{x}(r) + E(y \mid r, y \notin Y_0(r, c)) \le \tilde{x}$, i.e. that $E(y \mid r, y \notin Y_0(r, c)) \le 0$. Now,

$$E(y \mid r, y \notin Y_0(r, c)) = \frac{1}{A} \int_0^{c(r)} \int_{y \notin Y_0(c,r)} y d\Psi_y(y) dF(c)$$

$$= \frac{1}{A} \int_0^{c(r)} \left(\int_{-\infty}^{r - \tilde{x}(r) - d(c)} y d\Psi_y(y) + \int_{r - \tilde{x}(r) + d(c)}^{\infty} y d\Psi_y(y)\right) dF(c), \tag{4}$$

where

$$A = \int_0^{c(r)} \int_{y \notin Y_0(c,r)} d\Psi_y(y) dF(c) > 0.$$

Notice then that,

$$\begin{aligned}
&\int_{-\infty}^{r-\tilde{x}(r)-d(c)} y d\Psi_y(y) + \int_{r-\tilde{x}(r)+d(c)}^{\infty} y d\Psi_y(y) \\
= &\int_{-\infty}^{-(r-\tilde{x}(r))-d(c)} y d\Psi_y(y) + \int_{-(r-\tilde{x}(r))-d(c)}^{(r-\tilde{x}(r))-d(c)} ry d\Psi_y(y) + \int_{r-\tilde{x}(r)+d(c)}^{\infty} y d\Psi_y(y) \\
= &\int_{-(r-\tilde{x}(r))-d(c)}^{r-\tilde{x}(r)-d(c)} y d\Psi_y(y) \leq 0 \qquad (5)
\end{aligned}$$

where the inequality follows, since $r - \tilde{x}(r) = (\frac{1}{\lambda} - 1)r > 0$ by assumption, $d(c) > 0$, and hence the interval of integration is centered on a negative number while the normal distribution is symmetric about zero. Furthermore, the inequality is strict whenever $\sigma_y^2 > 0$ and $0.5 < \tilde{\lambda} < 1$. Consequently, the whole integral in (4) must be negative. When $r < 0$ the inequality in (5) is simply reversed proving the proposition. □

Proof. Proof of Corollary 1. It is enough to show that

$$\frac{1}{A'} \int_{-\infty}^{\infty} \int_0^{c(r)} \int_{y \notin Y_0(c,r)} r(x) y d\Psi_y(y) dF(c) d\Psi_x(x) \leq 0, \qquad (6)$$

where

$$A' = \int_{-\infty}^{\infty} \int_0^{c(r)} \int_{y \notin Y_0(c,r)} d\Psi_y(y) dF(c) d\Psi_x(x) > 0$$

and Ψ_x is the cumulative distribution function of a zero-mean standard normal distribution with variance equal to σ_x^2. It can be verified that multiplying the integrand in the proof of the previous proposition by r does not change inequality (5) when r is positive and reverses it when r is negative. Consequently, the inner double integral in (6) is always less than zero proving the corollary. □

Lemma 3. *$r(x)$ is weakly increasing.*

Proof. Proof. Let $\pi(r, x)$ denote the interim expected profits of the algorithm given a signal x and a report r. Recall that

$$\pi(r,x) = \int_{c(r)} \left(\int_{y \in Y_0(r,c)} \left((x + y - r)^2 + \sigma_z^2 \right) d\Psi_y(y) + \int_{y \notin Y_0(r,c)} (\tilde{x}(r) - x)^2 d\Psi_y(y) \right) dG(c) \\ + (1 - G(c(r)))x^2$$

Rewritting we have

$$\begin{aligned}\pi(r,x) &= \int_{c(r)} \left(\int_{y\in Y_0} \left((y-r)^2 + \sigma_z^2 + 2(y-r)x\right) d\Psi_y(y) + \int_{y\notin Y_0} \left(\tilde{x}(r)^2 - 2x\tilde{x}(r)\right) d\Psi_y(y) \right) dG(c) \\ +x^2 & \\ &= A(r) - B(r)x + x^2,\end{aligned}$$

for some non-negative functions $A(r)$ and $B(r)$. It follows that for every pair r, r', the set $X(r) := \{x \geq 0 : \pi(r,x) \geq \pi(r',x)\}$ is convex (and analogous for $x < 0$). This rules out the existence of a triple $x_0 < x_1 < x_2$ with $r(x_0) = r(x_2) \neq r(x_1)$. Hence, we can assume that for any x belonging to a decreasing segment of $r(x)$, (x_0, x_1), $\tilde{x}(r(x_0)) = x$. Hence,

$$B(r) = \int_{c(r)} \left(\int_{y\in Y_0} 2(r-y) d\Psi_y(y) + \int_{y\notin Y_0} 2r^{-1}(x) d\Psi_y(y) \right) > 0.$$

□

C Evidence on Adjustment Costs

Here we reproduce for ease of access some of the evidence already shown in Garcia et al. (2022) arguing that the hotel managers' behavior is consistent with them facing adjustment costs when changing prices. The hotels in that paper are the same as in our sample, we only changed their labels to retain the hotels' anonymity.

First, Figure 8 plots the relative frequency of price changes and recommendation changes for the two biggest hotels in the sample. The main takeaway from this figure is that the managers seem to have clear workday patterns that are not mirrored in the frequency with which the recommendations change. For example, the manager at Hotel 6 seems to concentrate on other tasks than pricing on Thursdays and Sundays, and does a lion's share of her pricing decisions on Tuesdays and Saturdays. This pattern suggests that the opportunity cost of time used on pricing is significant and varying over time. The pattern is consistent across all hotels in the sample as is evident from Table 9 which shows the share of all price updates done on each weekday for each hotel. As can be seen from the table, most hotels have at least one day on which they update next to no prices and often another day on which they do a large share of their updates.

Furthermore, Garcia et al. (2022) argue that Figure 12 strongly suggests that copying prices is less costly for the hotel manager than manually adjusting them. In that figure we see the distribution of the logarithm of total number of price updates separately for days

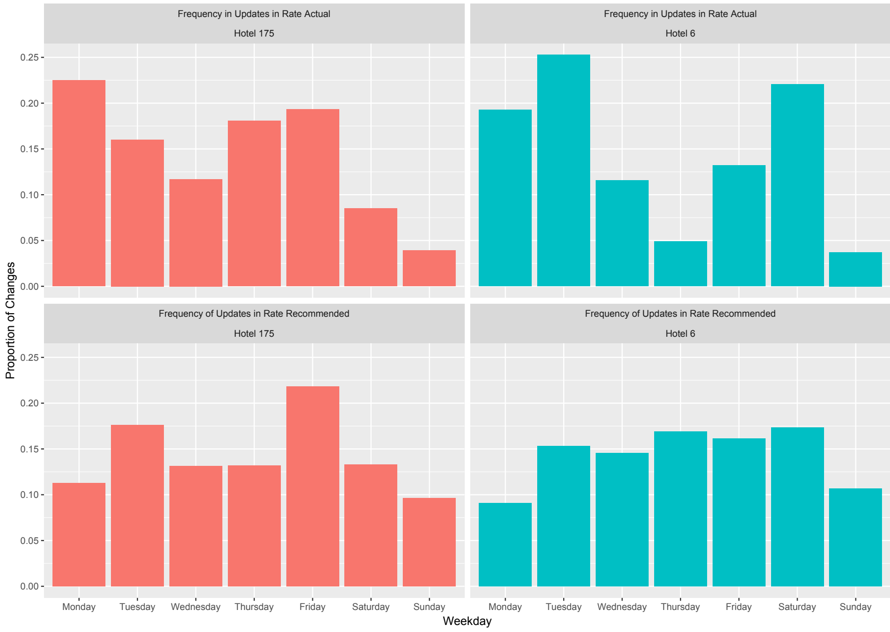

Figure 8: Frequency of Updates in Prices and Recommended Rates Across Weekdays for Hotel 6 and 175. Source: Garcia et al. (2022).

when the manager copies the recommendation for at least a full arrival week for at least one room type and for days when this does not happen. We see that on days when the manager copies recommendations she adjusts considerably more prices. This suggests that manually adjusting prices takes significantly more of the manager's time and effort. More evidence on the adjustment costs are provided in Garcia et al. (2022).

D Identification

In this section we show that a restricted version of the model is directly identified. In particular, we assume here that $\sigma_y = 0$. Because the manager obtains no information additional to r prior to paying c_2, $Corr(p, r \mid p \neq r) = \lambda$ and

$$E\left((p - E(p \mid r, p \neq r))^2 \mid r, p \neq r\right) = \sigma_z^2.$$

Similarly, $\lambda^2 \sigma_x^2 = Var(r)$. We need only show that with this information we can now identify the parameters of the cost functions. First, let q_1 denote the unconditional probability of a manual price adjustment. It follows that

$$\int G\left(\left(\frac{1-\lambda}{\lambda}\right)^2 x^2 + \sigma_z^2\right) d\Psi_x(x) = q_1,$$

Table 9: Distribution of Actual Rate Updates

Hotel ID	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday
6	0.19	**0.25**	0.11	**0.04**	0.12	0.22	0.04
10	0.19	**0.06**	**0.27**	0.09	0.07	0.16	0.17
11	0.17	**0.47**	0.02	0.04	0.18	0.12	**0.00**
23	0.14	0.14	0.19	0.12	**0.24**	0.10	**0.08**
30	0.22	0.10	0.17	**0.25**	0.16	0.06	**0.03**
131	0.14	0.15	**0.16**	**0.07**	0.16	0.17	0.16
175	**0.22**	0.16	0.12	0.19	0.19	0.08	**0.04**
192	0.12	0.16	**0.30**	**0.04**	0.14	0.11	0.13
208	0.07	0.31	0.07	0.08	**0.36**	**0.03**	0.06

Notes: Numbers in bold indicate the day with maximal or minimal density of actual rate updates for each hotel. Rate updates for each hotel sum to 1, rounding errors may apply. Data includes only products for which we observe $T \geq 100$ days before arrival. Source: Garcia et al. (2022).

where $G(x)$ is the distribution of c_2. Similarly, let q_2 denote the probability conditional on the recommendation exceeding some value $r_0 = \lambda x_0$. It follows that,

$$\frac{1}{1-\Psi_x(\lambda x_0)}\int_{\lambda x_0} G\left(\left(\frac{1-\lambda}{\lambda}\right)^2 x^2 + \sigma_z^2\right) d\Psi_x(x) = q_2.$$

Since $G(\cdot)$ is a two-parameter distribution (b_2, σ_c) these two moments pin it down. Finally, recall that $c(\lambda x)$ is the maximum cost such that a manager who observes a recommendation $r = \lambda x$ adjusts the price. The function $c(r)$ can now be computed in closed-form using the estimated parameters. In particular,

$$c(r) = \frac{(1-\lambda)^2 r^2 + \sigma_z^2}{b_1 + b_2} 1_{r<r_0} + \frac{(1-(1-\lambda)^2) r^2}{b_1} 1_{r>r_0},$$

with r_0 such that

$$\frac{(1-\lambda)^2 r_0^2 + \sigma_z^2}{b_1 + b_2} = \frac{(1-(1-\lambda)^2) r_0^2}{b_1}$$

It follows that

$$\int G\left(\frac{b_1 c\left(\frac{x}{\lambda}\right)}{b_2}\right) d\Psi_x(x) = q_0,$$

where q_0 is the unconditional probability of a price change and we use the fact that b_1 and b_2 are the respective scalers of the distribution.

Introducing $\sigma_y > 0$ represents a substantial increase in complexity but it is necessary to reconcile the model with the observed pricing behavior in the data. First, for some hotels

Table 10: Percentage changes in actions and losses relative to best reponse

	Percentile of F_x										
	0.50	0.55	0.60	0.65	0.70	0.75	0.8	0.85	0.90	0.95	0.99
Δ r	100%	1.0%	3.3%	-15.4%	2.5%	-2.6%	-4.9%	-0.8%	-14.0%	-16.4%	-19.1%
Δ loss	-0.00%	-0.00%	-0.02%	-0.01%	-0.00%	-0.00%	-0.00%	-0.02%	-0.01%	-0.19%	-0.52%

Notes: The table compares the algorithmic advisor's best responses and payoffs relative to actions and payoffs implied by the linear strategy. If $r^*(x;\lambda)$ is the best response given signal x and the hotel manager's actions, the first row reports $100\% \times \frac{r^*(x;\lambda)-\frac{x}{\lambda}}{r^*(x;\lambda)}$ where λ is the bias estimated in the main section. Because the linear recommendation is always zero at the 50th percentile ($0/\lambda = 0$), the percentage change when deviating to the best response will mechanically be $+/-100\%$. Similarly, if $p(r^*(x))$ is the random variable that represents the price realization given the hotel manager's strategy from the baseline model and the algorithmic advisor best responding to it, and if $p(x/\lambda)$ is the implemented price without a deviation, then the second row reports the reduction in expected losses for the algorithm's designer from best responding, i.e. $100\% \times \frac{\mathbb{E}\left[(p(r^*(x))-p^*)^2\right]-\mathbb{E}\left[(p(x/\lambda)-p^*)^2\right]}{\mathbb{E}\left[(p(r^*(x))-p^*)^2\right]}$.

the implied bias is larger than 1/2, meaning that, conditional on r, $p = 0$ is closer to the ideal price than $p = r$. This would then be inconsistent with a substantial fraction of prices that match the recommendation. Second, the empirical distribution of $p - E(p \mid p \neq r)$ is double-peaked and has a valley around zero. This suggests that managers are less likely to change the price manually whenever some privately observed shock is small in magnitude, which is precisely what we capture with the variable Y.

E Gains from Best Responding

To gauge the restrictiveness of limiting the algorithm's strategy space to linear strategies we simulate the loss for the algorithm's designer, were she to privately deviate to her non-linear best response, assuming that the hotel manager still believes that the algorithm is using its linear strategy. We then compare this to the loss from the linear strategy. The expected payoffs are calculated as averages of payoff realizations over 10000 draws from the distribution of signals which we estimated for the pooled sample in the main text (see the first row in Table 6). We consider best responses to revenue manger's signal realizations for 10 evenly split percentiles starting from the 50th percentile. The results are reported in Table 10.

Notice that due to symmetric signal distributions the payoff losses for percentiles below the 50th percentile will be symmetric to the ones presented here. The results are calculated as a percentage of the best response outcome. We also estimate the results for the 99th percentile because high signal realizations are the ones where the algorithm already has a very high chance of inducing the hotel manager to revise her price even without exaggerating its recommendation.

For most signal realizations the difference in the advisor's payoff when best responding compared to when playing the linear strategy is negligible, see first row in Table 10. For all but the two highest percentiles of signals in the table, best responding reduces the advisor's losses by at most 0.02%. As mentioned above, for the higher signal realizations the advisor would like to reduce her lying by a significant margin but even this reduction will increase her payoff by only 0.19% at the 95th percentile of signals and by 0.52% in the 99th percentile. We conclude that the restriction to linear strategies does not seem to generate significant deviation incentives and hence is likely to have a negligible quantitative impact on the results.

F Additional Results for Counterfactuals

In the following, we provide some additional information on our counterfactuals. Figure 9 shows the estimated equilibrium bias when we vary the recommendation algorithm's share of the policy-relevant information similar to Figure 7 in the main text. Specifically, Figure 9 shows an initial increase in the recommendation bias followed by a leveling out and even a slight decrease in the bias when the recommendation algorithm holds most of the price information.

In Figure 10 we compare delegation against the current setup if we square the standard deviation of the normal distribution of manager's adjustment costs, σ_c.[34] The figure illustrates how the current environment can be superior to delegation when the hotel manager holds enough information and her price-adjustment costs are often sufficiently low to allow for manual adjustments.

G Asymmetry of Price Responses

In this section we present evidence that prices respond relatively symmetrically to both increases and decreases in the algorithmic recommendation. As a first piece of evidence, we show estimates from a Cox proportional hazards regression where the "failure" event is a

[34] This is practically the same as doubling the standard deviation, since $\sigma_c = 2.28$ and hence $\sigma_c^2 = 5.20$.

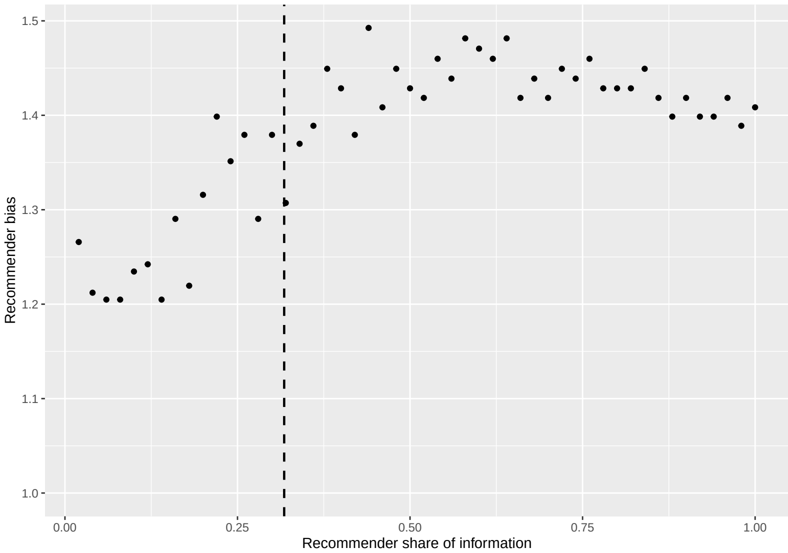

Figure 9: Estimated equilibrium bias as a function of the share of information held by the recommendation algorithm.

change in the price and we allow the probability of a price change to depend differently on positive and negative changes in the recommendation. We estimate the model allowing for different baseline hazard rates for different hotel-room type-arrival month combinations. More precisely, let T_{hrka} be the random next update time since the $(k-1)$th previous update time at which the price for the room type r in hotel h is being changed for arrival date a in month m. We estimate the model

$$\mathbb{P}(T_{hrka} \geq t) = \exp\left(-\int_0^t \lambda_{hrm}(t) \exp\left(\beta_1 \Delta^r_{hrkat+} * 100 + \beta_2 \Delta^r_{hrkat-} * 100 + \beta_3 Days\ ahead_{hrkat}\right) dt\right),$$

where λ_{hrm} is an arbitrary measurable baseline hazard rate that is allowed to be different for different combinations of h, r and m, the variable Δ^r_{hrkat+} is the absolute change in the logarithm of recommendations since the last price update if the recommendation increased and zero otherwise, Δ^r_{hrkat-} is the same when the change in the recommendation was negative and $Days\ ahead_{hrkat}$ counts how many days before the arrival date a the tth day after the kth update was.[35] We multiply all of the recommendation changes by 100, as typical recommendation changes are about 5% The β_1, β_2 and β_3 are regression parameters. We run this model excluding 0.5% of the largest changes in recommendations, as these changes never induce a price change, typically are followed by a reverse correction in the recommendation

[35] The results are quantitatively very similar if we include a third order polynomial of *Days ahead* or exclude it completely.

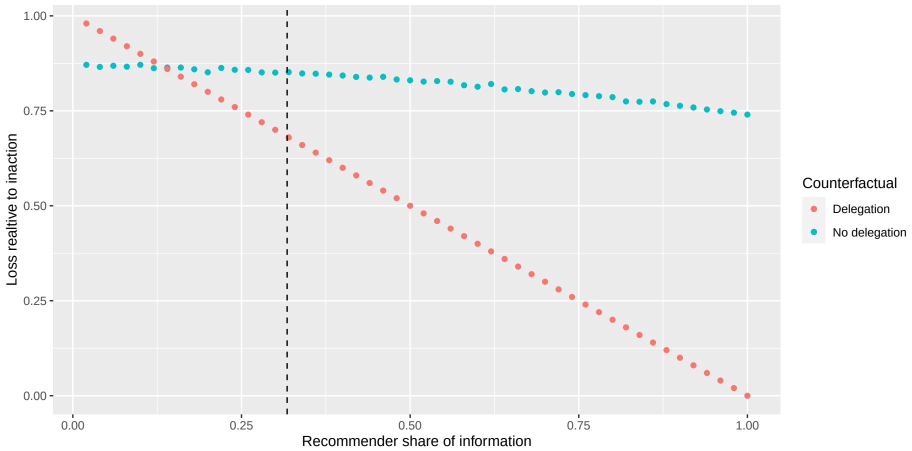

Figure 10: Counterfactual Information Shares and Hotel Manager Payoffs with Doubled Variance of Adjustment Costs

and hence seem to be considered as mistakes by both the revenue manager and the hotel manager.[36] Table 11 reports hazard ratios implied by a unit change in a covariate (which in the Cox model reduce to $\exp(\beta_i)$ for covariate i) and their 95% confidence intervals for the different changes in recommendation holding the other covariates constant.

As can be seen from the table, positive and negative price changes in the recommended price have a remarkably similar and fairly large impact on the probability of updating. A way to interpret the estimates is that a 1% increase or decrease in the price results in approximately a 16% increase in the hazard rate of a price update by the hotel. Similarly, an extra day ahead of arrival reduces the hazard rate of a price update by approximately 1%. In other words, both recommendation increases and decreases induce the hotel manager to update their prices and do so in a remarkably similar fashion.

As a second piece of evidence, figure 11 shows the relationship between log changes in the recommendation and log changes in the price conditional on a manual price change by the manager. The Figure shows also a linear fit of the data that allows for a discontinuity at zero and a different slope for positive and negative recommendation changes. As can be seen from the figure, there is next to no discontinuity and the slopes for increases and decreases are remarkably similar.

[36] Despite being only about 9000 of the total 2 million prices, as extreme observations, they have a large impact on the Cox estimates and including them would lead us to conclude that large changes in the magnitude of the recommendation *reduce* rather than increase the probability of a price update.

Table 11: Hazard ratios from Cox proportional hazards model for the probability of a price change as a function of the magnitude of the change in the recommendation

$\Delta^r_{hrkt+} * 100$	$\Delta^r_{hrkt-} * 100$	*Days ahead*
1.165	1.161	0.989
$[1.163, 1.1676]$	$[1.158, 1.1644]$	$[0.989, 0.9894]$

Notes: The first row reports hazard ratios and the second row reports 95% confidence intervals for those hazard ratios. The estimation sample excludes recommendation changes that are higher than 0.25 logpoints in absolute value, this corresponds to approximately 0.5% of the total price observations. N=2,008,202.

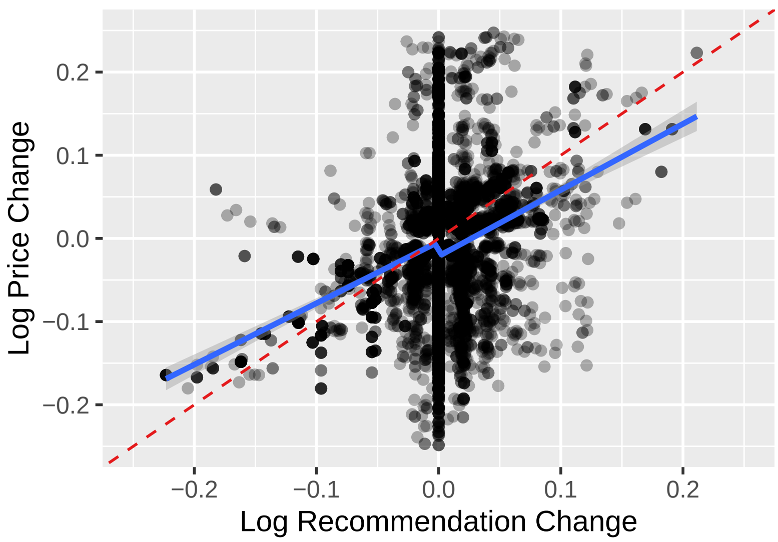

Figure 11: Model Fit: Recommendations and Prices

Notes: Each point represents a price update that does not match the recommendation. The horizontal axis is the log change in the recommendation and vertical axis is the log change in the price. The blue line shows a linear fit of the data with a 95% confidence interval allowing for a discontinuity at zero and a different slope on both sides of the origin.

H Alternative Price-Adjustment Models

In the following, we discuss alternatives to the price-updating model presented in the main text. We define a recommendation $r(x)$ to be unbiased if $E[p^* \mid x] = r$.

- Suppose that p^* is normally distributed with mean zero and standard deviation σ. Furthermore, assume that the manager observes signal y and the algorithmic advisor observes signal x, which conditional on p^*, are independent and both normally distributed with mean p^* and standard deviation σ_i, $i \in \{x, y\}$. The unbiased recommendation is $r(x) = \mathbb{E}[p^*] = \frac{\sigma^2}{\sigma^2+\sigma_x^2}x$ and p^* conditional on $r(x)$ is distributed normally with mean equal to $r(x)$. That is, an unbiased algorithmic advisor does not "naively" report x but instead *deflates* her signal. Notice that "naively" sending one's signal without deflating it is highly costly in our setting because the hotel manager often copies the recommendation. If the advisor truly shows a low level of sophistication and passes on its signal the model would correspond to the alternative specification in section 6 where λ is a primitive of the data and not a choice parameter. The counterfactual results for this model remain qualitatively similar to the original ones due to the high correlation between the bias estimated in the baseline model and the bias that would result from this low level of sophistication (see Figure 6).

- Consider the model above but assume that the hotel manager's signal is fully informative and she copies the recommendation when the difference between the recommendation and the truth does not warrant a manual adjustment cost and otherwise sets the price optimally. In other words, assume that $p = r$ if $p^* \in (r - c, r + c)$ for some $c > 0$ and $p = p^*$ otherwise. Under the hypothesis that the recommendation is unbiased

  $$E[p^* \mid p \neq r, r] = E[p^* \mid p = r, r] := r.$$

  In our data $E[p^* \mid p \neq r] = \gamma r$ for some $\gamma < 1$ which contradicts the equality above.

- Truth-noise model: Suppose that the information held by the algorithm is strictly worse than that of the manager. In particular, $x = y = p^*$ with probability q and otherwise x is an imperfect predictor of $y = p^*$. In particular, assume that, $E[p^* \mid x, x \neq p^*] = \gamma x$. In this case, we have that the recommendation is unbiased only if $r = (q + (1 - q)\gamma)x$ in which case

  $$E[p^* \mid p \neq r] = \frac{q(1-\gamma)}{q + (1-q)\gamma} r.$$

A prediction of this model is that, immediately upon observing $p \neq r$, the algorithm's recommendation should change to $r' = p$. In the data, we observe the algorithm updating immediately after a price change that does not match the recommendation with 83% probability, see Figure 13, but only 13% of these updates result in $r' = p$ and $E(r' \mid p, p < r) > p$, i.e. the recommendation does not fully react to the change in price as shown in Figure 14. Together these suggest that there is persistent "disagreement" between the algorithm's designer and the hotel manager.

- Intrinsic Attention: Our model assumes that recommendations drive attention allocation. An alternative hypothesis is that the manager devotes attention to those products she obtained some information about and uses the recommendation as a confirmation/shortcut. That is, the manager first observes y and decides whether to pay attention and if so then observes r, choosing whether to accept or reject the recommendation. We contend that this model is implausible for a number of reasons. First, if the manager only puts attention when observing extreme values of y (because of attention and adjustment costs), then the expectation of the difference between the recommendation and y conditional on the manager putting attention to a price would be large (even if we allow for x to be correlated with y) resulting in a low likelihood of copying the recommendation. To match this moment, it then should be the case that r is very close to y almost always, rendering the information held by the manager useless. Instead, our timing assumes the manager devotes attention when r draws an extreme value and the manager accepts if y is relatively small which occurs much more often. Second, we observe hotel managers accepting hundreds of recommendations in one day while not changing a single price manually, see Figure 12.

- Revenue vs. profits: Since the revenue management company is benchmarked on revenue but the hotel manager *should* care about profits, there could be a *directional* disagreement between them based on this difference in payoffs. Indeed, theoretically, the revenue-maximizing price is always lower than the profit-maximizing price. Huang (2022) shows that Airbnb's recommendations to hosts are downward biased, consistent with the platform's preference for revenue-maximization. By contrast, we do not observe a significant directional bias in our data: the recommendation algorithm exaggerates both price hikes and drops by approximately the same amount. This is illustrated in Figure 11 which shows the dependence of manual updates on the recommendation. The figure also plots a linear fit of the data which allows for a discontinuity at zero and dif-

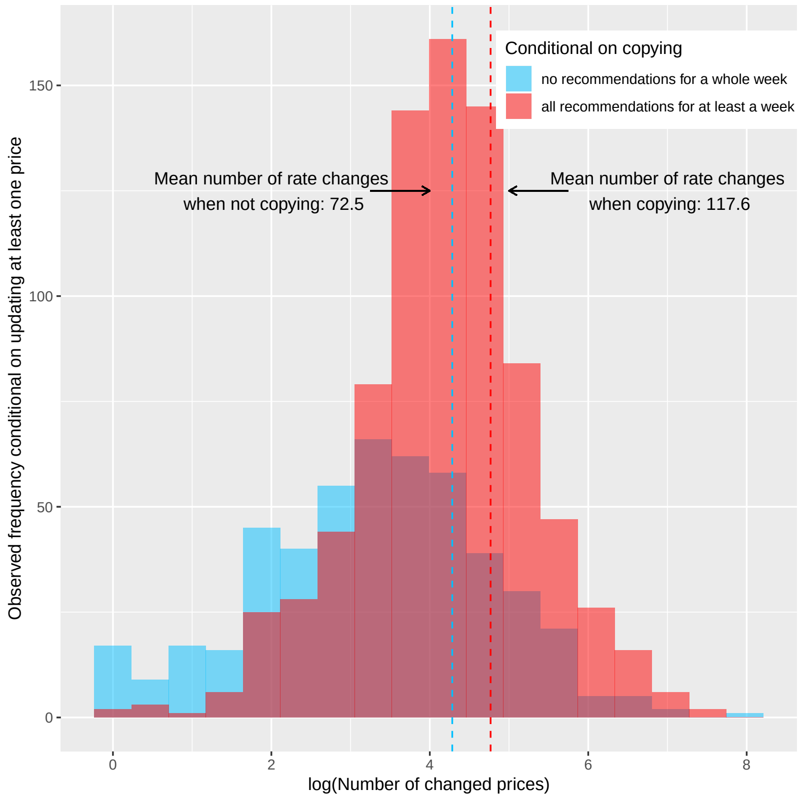

Figure 12: Number of Rate Updates (in Log) Conditional on Copying and Not Copying Recommended Rates. Source: Garcia et al. (2022).

ferent slopes on both sides of the origin. As can be seen from the figure, there is only a marginal discontinuity at zero and the slopes on both sides of the origin are remarkably similar.

I Details on Estimation and Robustness

Our main estimation routine proceeds as follows:

1. Fix a starting λ and minimize the quadratic distance of the model implied moments to the empirical moments standardized by the empirical moment.[37]

2. Check if local deviations in λ at the found parameter values are beneficial. If yes, update λ to the direction of the beneficial deviation and return to step 1. If not, end

[37] I.e. If the model implied moments for a parameter vector θ are $\mu(\theta)$ and the empirical moments are given by a vector m, we solve $\min_\theta \{\mu(\theta)' I_{m^{-1}} \mu(\theta)\}$, where $I_{m^{-1}}$ is a diagonal matrix where the ith diagonal element is $\frac{1}{m_i}$.

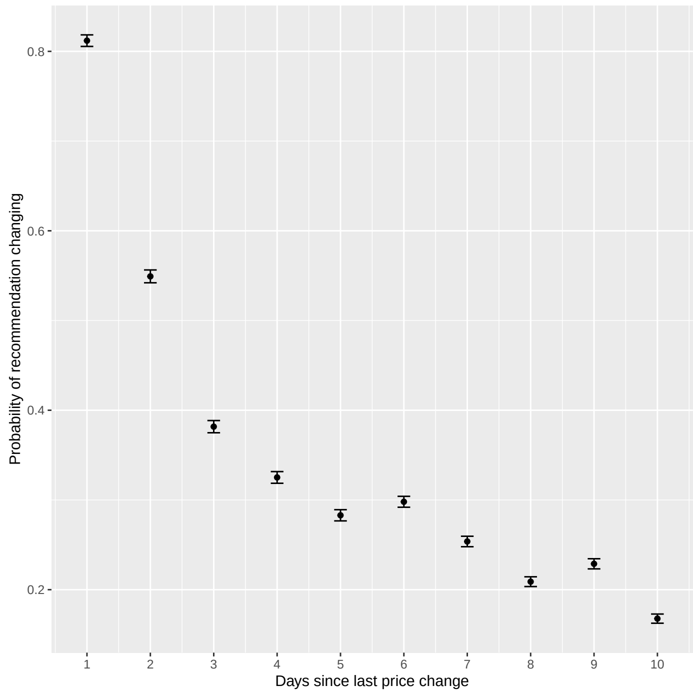

Figure 13: Probability of the recommendation changing after the price manually changed to not equal the recommendation. Ticks represent the 95% confidence intervals.

the routine.

Because this routine is quite sensitive to the intial values of λ, the current starting values are the ones generating the lowest quadratic scaled distance above that we have been able to find after starting the routine from dozens of different values for the bias.[38] However, irrespective of the estimation procedure, the counterfactuals remain very stable. To illustrate this we report in Table 12 the results of the counterfactuals if we treat λ as a parameter and estimate it simply using SMM like the other parameters in the first step of our baseline routine but do not require it to imply a perfect Bayesian equilibrium. The resulting counterfactual losses are remarkably close to the ones estimated with the original estimation routine.

[38] Due to computational load we have not been able to implement of proper grid search for the loss minimizing set of parameter values.

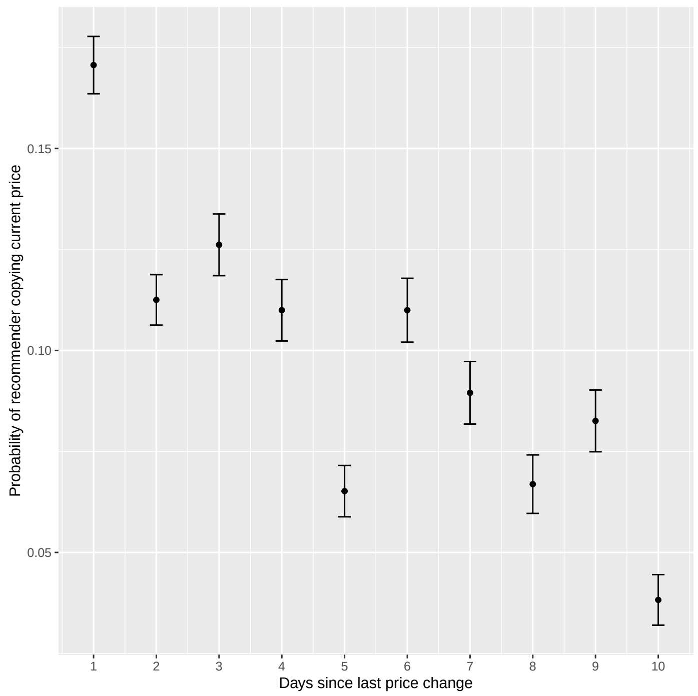

Figure 14: Probability of the algorithm copying the current price after the price manually changed to not equal the recommendation. Ticks represent the 95% confidence intervals.

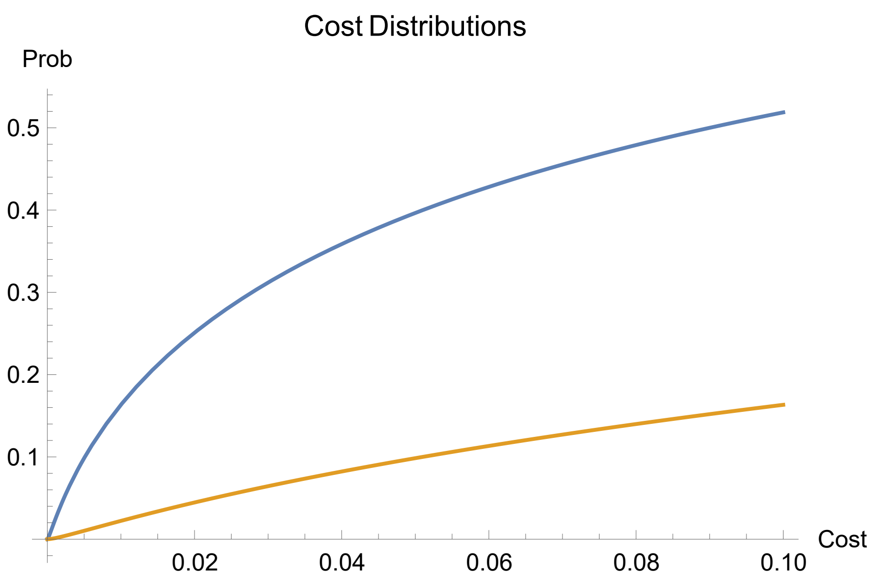

Figure 15: Cumulative distribution functions of cost c_1 for copying recommendations (in blue) and cost c_2 for manual price adjustments (in yellow).

Table 12: Counterfactuals from Full SMM

Hotel	Benchmark	Profit Loss	Delegation	Biased
A	0.996	0.992	0.975	0.982
B	0.985	0.971	0.820	0.822
C	0.982	0.966	0.726	0.735
D	0.992	0.985	0.934	0.950
E	0.983	0.967	0.756	0.773
F	0.992	0.986	0.710	0.712
G	0.980	0.966	0.639	0.667
H	0.990	0.978	0.844	0.864
I	0.987	0.973	0.817	0.823

Notes: The value in the first column (Benchmark) corresponds to the welfare loss in the status quo relative to the welfare loss under complete inaction. The value in the second column (Profit Loss) is the implied accounting profit loss, disregarding adjustment costs, relative to inaction. The third column (Delegation) represents the welfare loss in the counterfactual exercise of full delegation to the algorithm, again relative to inaction. The fourth column (Biased) describes the expected welfare loss from a counterfactual where the decision is delegated to the algorithm which continues to produce biased recommendations relative to complete inaction. The last column (No Rec) shows the expected welfare loss relative to inaction if the hotel manager has no access to the recommendation or copying it.