

Strategic Responses to Algorithmic Recommendations: Evidence from Hotel Pricing*

Daniel Garcia[⊕] Juha Tolvanen[⊖] Alexander K. Wagner[⊗]

July 23, 2022

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 sub-optimal pricing by human decision makers. We quantify the losses from this strategic bias in recommendations using a simple structural model and estimate the potential benefits that would result from a shift to fully automated algorithmic pricing.

JEL-classification: D22, D83, L13

Keywords: Advice; Algorithmic Recommendations; Delegation.

* We would like to thank seminar audiences at TSE, CEU, UNLP and the 2021 CESifo Economics of Digitization workshop for helpful comments and suggestions. Garcia gratefully acknowledges support from FWF FG Pricing in Imperfectly Competitive Markets. The computational results presented have been achieved in part using the Vienna Scientific Cluster (VSC).

[⊕] Department of Economics, University of Vienna, Oskar-Morgenstern-Platz 1, A-1090 Vienna, Austria. Email: daniel.garcia@univie.ac.at

[⊖] Department of Economics, University of Vienna, Oskar-Morgenstern-Platz 1, A-1090 Vienna, Austria. Email: juha.tolvanen@univie.ac.at

[⊗] Department of Economics, University of Salzburg, Mönchsberg 2A, A-5020 Salzburg, Austria, and VCEE, University of Vienna, Austria. email: alexanderk.wagner@plus.ac.at

1 Introduction

Organizations rely increasingly on prediction algorithms for decision making, with application 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 in online platforms and traditional markets (Chen et al., 2016; Brown and MacKay, forthcoming; Assad et al., 2020; Calvano et al., 2020; Garcia et al., 2022). In many of these applications machines augment human decisions by acting as advisors to human managers who retain the final decision right. It is theoretically 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 The empirical evidence on how algorithmic advice and human decision making work together in *strategic situations* is however extremely limited (Athey et al., 2020; Cowgill and Stevenson, 2020).

In this paper, we study the strategic 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 model of advice to identify a novel source of misalignment between the algorithmic advisor and the human decision maker originating from price-adjustment costs humans face. We show that this conflict of interest can lead to substantially distorted recommendations by the algorithm and suboptimal pricing decisions by human managers in equilibrium. 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, see Section 3 for details. Using a simple structural model, we quantify the resulting losses from mispricing. Our main counterfactual experiment shows that full delegation to the recommendation algorithm significantly outperforms the current organizational setting in which human managers set prices for most hotels in our sample.

In our setting, both the revenue management company and the hotel manager benefit from maximizing the hotel’s profits.² The main source of conflict in this paper arises be-

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.

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

cause the designer of the algorithm, in contrast to the manager, does not incur adjustment costs when (automatically) changing prices and hence would like to implement price updates more frequently.³ Because larger deviations from the optimal price are more costly to the hotel manager, the algorithm’s recommendation *exaggerates* the change in the optimal price to induce the manager to update prices faster. This, 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 whenever the more exaggerated recommendation prompts more frequent price updates from the manager.

We believe that this type of conflict of interest is likely to arise in most settings where the decision maker has an easy default or status quo option (e.g. ‘keep current price’) and deviating from that option (e.g. ‘update price’) incurs some form of adjustment costs.⁴ An algorithmic recommendation for an alternative action will have two beneficial effects in addition to the direct benefit of providing better or more aggregated information. First, it simplifies the decision maker’s task by offering an easily selectable alternative. Second, its recommendations may induce the decision maker to more carefully consider other alternatives than the status quo or the algorithm’s recommendation leading to better decisions. Both effects are likely to also benefit the algorithm’s designer. However, the algorithm’s designer prefers more frequent decisions than what is privately optimal for the decision maker, and is hence tempted to bias its recommendations if it can induce faster adjustments away from the default option.

We build a stylized empirical model of advice that captures the key features of the price-setting interaction of the algorithmic advisor and the manager. In this model, we assume that human managers display limited attention and incur price-adjustment costs.⁵ The manager observes costlessly the price and the 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 devote costly attention to it. If she chooses the latter,

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 For example, a doctor can easily default to the most common cause of the most prominent symptoms when making a diagnosis, parole boards can keep the inmate locked in, and a store manager can choose to not restock their inventory.

⁵ Different interpretations of adjustment costs are discussed in Section 5. There is also a substantial literature in macroeconomics that studies adjustment and information processing costs in price setting (e.g. [Alvarez et al., 2011](#)).

she receives an informative signal about the optimal price and then decides whether to copy the recommended price. If she does not copy, she pays an additional adjustment cost to learn more information and gets to update prices freely. The assumed timing and cost structure for adjusting prices is reflected by the pricing interface managers use: accepting 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.

The model of strategic advice with costly price adjustments successfully matches the key empirical observations of the relationship between algorithmic recommendations and realized prices. First, we show in Section 4 that actual price updates are relatively rare compared to updates in price recommendations. Over the complete booking horizon, human managers update prices for a particular product about once a month whereas algorithmic recommendations are updated four times more frequently. This difference in updating frequencies, together with rarely observed small price changes (within 1.5% of the current price) 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 change in the recommendation. 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 after a large than a small 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.⁶ 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.⁷ This pattern is otherwise hard to reconcile with standard models of advice as we discuss in Section 2.

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 to 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

⁶ Correlation between the two adjustment costs is highly plausible because both represent an opportunity cost of the manager's time.

⁷ 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.

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 information in case of manually updating prices further decreases the empirically observed pass-through rate of recommended prices.⁸ The reason is that manual price changes are more profitable in situations in which the manager’s information strongly contradicts the recommender’s signal. Taken together, the model is rich enough to generate the, at first sight, counterintuitive empirical pattern of (i) a high unconditional copy rate as well as (ii) a considerably dampened pass-through of recommendations in case the manager decides to adjust prices manually.

A primary ingredient of our model presented in Section 5 is the perceived bias in the recommendation. 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 acquire it, implying that hotel managers must be facing substantial adjustment costs. This leads to inaccurate decisions: the standard deviation of the difference between the optimal price and the actual price is larger than the standard deviation of the price

⁸ In other words, our model predicts that the actual bias in recommendations is smaller than the gap between recommended changes and manual price changes.

observed in the data.

In section 7, we finally address whether full delegation to the algorithm would lead to better pricing decisions. Delegation to the algorithm has some obvious benefits. It brings about instantaneous and costless decision making and aligns preferences of the algorithmic advisor and the human manager. A potential cost of delegation however is that a hotel manager’s private information is not available as an input for decision making. Our counterfactual experiment shows that delegation to the algorithm would reduce expected losses by 8 to 35 percent in the sample. Decomposing these gains shows that the majority of them 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%).

We discuss the implications and possible other applications of our results on the strategic interaction between algorithmic advice and human decision makers in Section 8.

2 Related Literature

We contribute to a thriving literature studying the interaction between algorithmic advice and human decisions focusing on an economically relevant application: managerial pricing. There exists a large theoretical literature on strategic advice in economic organizations (e.g. [Kamenica and Gentzkow, 2011](#); [Sobel, 2013](#); [Kamenica, 2019](#)).⁹ In the canonical setup, 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 attention as the source of disagreement. We believe this is a relevant consideration in many organizations, whereby generalist managers rely on information from several specialized experts. In the context of the interaction between human decision makers and machine advisors, this is almost certainly guaranteed to occur, as the adjustment cost of the machine is infinitesimal compared with that 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 algorithm.

⁹ [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.

A similar tension arises in [Kartik et al. \(2007\)](#), which introduces a cheap talk model 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 in 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 recommendations (akin to delegation, or rubber-stamping) *more* often, precisely whenever the recommended price is closer to the current price. The crucial insight is that also rubber-stamping requires, in contrast to fully automatic decision making, at least some attention from the human principal.¹⁰

[Bloedel and Segal \(2020\)](#) study a persuasion model where the principal is subject to rational inattention.¹¹ Just like 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, note that we do not observe pooling or bunching in the data. On average, prices set manually by the hotel manager increase continuously in the recommendation.

More broadly, we contribute to the empirical literature on strategic communication in organizations.¹² We are aware of two papers that use equilibrium analysis to iden-

¹⁰ 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.

¹¹ There is a very large literature on rational inattention, with some important applications to organizations; see [Maćkowiak et al. \(forthcoming\)](#), esp. Section 3.3).

¹² 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

tify 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).¹³ An exception is [Bundorf et al. \(2019\)](#) who study 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. 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). Our data contains bookings and prices for each hotel over a period of about 14 months. We observe for each room and each possible arrival day the flow of

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.

¹³ Prominent examples in this stream of literature are [Hoffman et al. \(2018\)](#), [Kleinberg et al. \(2018\)](#), [Ribers and Ullrich \(2019\)](#), [Chan et al. \(2022\)](#) and [Currie and MacLeod \(2017\)](#).

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 the actual price as the main instrument for price optimization.

A key input for our analysis is the recommended price. 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.¹⁴

The hotel manager decides every day whether to use the revenue management system to change 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 B 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 is less time-consuming than manually adjusting prices. Further evidence and discussion of the details can be found in [Garcia et al. \(2022\)](#).¹⁵

¹⁴ 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.

¹⁵ 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. Consequently, we think that the two settings are fairly comparable in terms of the

Table 1: Price Updates and Recommendations

	<i>Main Sample</i>			<i>Full Sample</i>		
	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 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 it matching the recommendation exactly (Copy Rec).

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 sample to include only those observations whereby the initial price matched the recommended price. This selection allows us to interpret differences between the final 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 sample includes approximately 34% of observations and 58% of the price updates.

opportunity and complexity costs of adjusting prices.

4 Stylized Facts: Recommendations and Prices

In this section, we present a number of novel empirical facts about the relation between price recommendations by the algorithm and price updates by the human decision maker. For this descriptive analysis observations of all hotels in our sample are pooled, but our findings also hold qualitatively for each hotel. The following key observations inform the choice of our model of price adjustments we present in Section 5.

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

The first observation is that hotel managers adjust prices only infrequently, on average, once every 35 days; with considerable heterogeneity across hotels as shown in Table 1. Because algorithmic recommendations change much more frequently, once every seven days, the difference in updating frequencies lead 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.*

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




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.

price updates 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,

¹⁶ This updating feature manifests itself in different shapes across different models. 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

Table 2: Price Update Probability

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

Notes: Fixed-effects regressions; 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: *** $p < 0.001$

we will account for the hotel manager’s distinctive updating behavior 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. Including various controls, such as room type-arrival week fixed effects and a polynomial of the days before arrival, leads only to a modest

assumes control when stakes are high, thus reducing the influence of the agent.




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.

increase in the estimated coefficient (73%). It follows that hotel managers must believe that the pricing algorithm exaggerates the optimal price change on average.

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.




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.

5 Model

We introduce in this section a stylized model of price adjustments with algorithmic recommendations and costly information acquisition of the hotel manager that can rationalize the empirical facts about the relationship between recommendations and pricing decisions. To ease the mapping of the model to the data, we 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 $i = x, y, z$. In our empirical specification, x , y , and z are assumed to be normally distributed, but our

¹⁷ This profit function can be micro-founded assuming a log-linear demand and semi-elasticity η .

Table 3: Price Update Copy Rates

Rec Change	<i>Copying Probability</i>			
	0.133*** (0.012)	0.141*** (0.011)	0.151*** (0.011)	0.151*** (0.011)
Hotel \times Date FE	No	Yes	No	No
Room \times Date FE	No	No	Yes	No
Room \times Date \times 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: *** $p < 0.001$

theoretical results do not depend on this feature.

Figure 4 illustrates the hotel manager’s sequential information-acquisition process and pricing decision. 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 but is not known to the hotel manager. The manager only learns about x by observing the algorithm’s price recommendation $r(x)$.

The manager then decides whether to allocate attention to adjusting prices. If she does not, she maintains the current price, $p = 0$, which incurs no cost. In case she does, she acquires information y for the simple adjustment cost c_1 . We think of y as information which is directly available from the hotel manager’s knowledge of events about a particular product, e.g. the hotel’s chef called in sick. Given the hotel manager decides to learn y , she can choose to (i) update the price by copying the recommendation resulting in $p = r$, or (ii) to acquire additional information z for the complex adjustment cost c_2 . Only in the case of learning z for cost c_2 , the manager can update the price freely, such that, $p = E(p^* | r, y, z) = E(x | r) + y + z$.

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 $b_i > 0$.

Table 4: Pass-Through Rates of Recommendation

	<i>Change in actual price</i>			
	All	Manually Updated		
Rec Change	0.974*** (0.002)	0.725*** (0.005)	0.733*** (0.006)	0.738*** (0.006)
Days ahead Polynomial	No	No	Yes	Yes
Room \times 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 \times manual. *Room* is the room type and *Month* refers to the arrival month. Significance levels: *** $p < 0.001$

Although we are agnostic about exact psychological nature of the two adjustment costs, simple and complex, the cost structure lends itself to two different interpretations. Parameter c can be understood as the opportunity cost of a unit of time for the manager and b_i as the time required to learn the information. Alternatively, c can be thought of as cognitive costs with b_i measuring the complexity of the task.¹⁸ Intuitively, it seems easier to decide whether the recommendation is satisfactory than to fully determine the optimal price manually. This structure is reinforced by the pricing interface, which allows accepting recommendations with a single click, while freely adjusting prices requires the hotel manager to access an additional screen and enter each price manually. Nevertheless, neither our theoretical nor empirical model imposes any assumption on which cost is bigger.¹⁹

¹⁸ This interpretation is consistent with decision models of limited attention (e.g. [Dean et al., 2017](#)), which show 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 any price p).

¹⁹ Another interpretation for these 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 the costs stems from managerial inattention.




Figure 4: Information acquisition and pricing of hotel manager

A key ingredient of the model is the algorithmic price recommendation. We assume that the algorithm’s designer cares about the hotel’s profits, $\Pi = \Pi_0 - \eta(p - p^*)^2$. 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 $\lambda > 0$.²⁰

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 algorithmic 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.

²⁰ This restriction on the algorithm’s strategy space makes both the theoretical model and the empirical model more tractable. Furthermore, we illustrate in Appendix D that the loss from restricting oneself to a linear reporting strategy is negligible for the algorithm’s designer.

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 A. 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 long 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,

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

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.

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 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, $\mu(0) > 0$, $\lim_{r \rightarrow 0} \mu(r) = 0$, 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

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

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*

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

and

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

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.*

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

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 model presented 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 then 6 structural parameters: three of them govern the informational environment $(\sigma_x, \sigma_y, \sigma_z)$ and three that correspond to the distribution of shocks

Table 5: Targets for Pooled Data

Moment	Data	Model
$\sqrt{Var(p Update)}$	0.074	0.074
$\sqrt{Var(r Update)}$	0.068	0.067
$\sqrt{Var(p - r Update)}$	0.035	0.035
$\sqrt{E(p \cdot r Update)}$	0.068	0.066
Pr(Copy Update)	0.840	0.842
Pr(Copy Update, Large Rec)	0.947	0.951
Pr(Update)	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 copy rate, both unconditionally and conditional on the recommendation change exceeding 10% (Large Rec). The last row reports the unconditional update rate.

(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. To estimate these parameters, we use a method of simulated moments, 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. Four of those targets, see row 1 to 4 in Table 5, depend directly on the joint distribution of recommendation and price updates. In addition, we match the likelihood of the price matching the recommendation, both unconditionally and conditionally on the recommendation change exceeding 10% as well as the average update rate.

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 local deviations are not beneficial. 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.

To discuss identification, it is instructive to consider a special case of the model in which $\sigma_y = 0$.²¹ Because there is no selection into updating based on payoff-relevant information, 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 .

We first run the routine on the pooled dataset. Results are summarized in the first row of Table 6, including bootstrapped standard errors. We find that the private information of hotel managers accounts for less than 20% in the total variance of the optimal price. This means that managers’ private information is at least five times as valuable as that of the algorithm. Unfortunately, accessing this information requires substantial effort by hotel managers. We estimate a mean adjustment cost for c_1 of approximately 0.2, with a standard deviation of 1.10.²² The cost of acquiring further information is estimated to be an order of magnitude larger, with mean 3.06 and standard deviation of 4.24. As a consequence, there is considerable dispersion between actual prices and counterfactual optimal prices. This disparity also implies that adjustment costs do indeed reflect costly

²¹ A proof of identification for this case, as well as a discussion of why this assumption is unlikely to hold, is provided in Appendix C.

²² Given our parametrization, the mean of the distribution is $\exp(\ln b_1 + \sigma_c^2/2)$.




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.

managerial attention rather than fear of consumer backlash (Rotemberg, 2005).

Our estimates also suggest a significant bias in recommendations, with $\lambda = 0.83$. Because most price updates match the recommendation exactly, this bias implies biased (sub-optimal) prices. Nevertheless, the welfare impact of the recommendation bias is ameliorated by the fact that the manager selects into the decision to match the recommendation, see Proposition 2.

Table 5 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 (see Figure 6).

On average, we find that the current regime does a relatively poor job at exploiting the available pricing information and is able to reduce losses from mispricing by less than 4%. This is, however, not very surprising because managers only rarely update prices and, in case they do, they often copy the recommendation.

We then run our estimation routine separately for each hotel. Results are summarized for hotel A to H in Table 6. Most hotels are well-represented by the pooled data. The variance of 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 available to the manager ($\sigma_y + \sigma_z$). In order to perform counterfactuals in the next section, we compute the expected loss in profit for each hotel with the current institutional setting, relative to the profit loss they would experience if they never updated their prices. 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

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

Notice that lower values correspond to more efficient outcomes, with $w_i = 0$ being the first-best outcome.

7 Counterfactual Delegation

In light of the estimation results presented in Section 6 it appears that a majority of hotels would benefit from fully delegating decision making to the algorithm. Delegation of pricing decisions to a fully automated algorithmic system has a number of advantages. It completely eliminates information processing and menu costs of hotel managers and thereby also removes delay in decision making. It should also lead to truthful recommendations because we identified the inertia of the hotel manager as a source of conflict with the algorithm in our application. A potential downside of the full delegation is that the algorithm does not have access to the realizations of signals y and z , which are even more valuable

than signal x according to our data. The key insight, however, is that the manager’s inertia reveals that these informative signals come at significant costs, which greatly reduces their value for decision making.

Our parameter estimates directly allows us to compute the residual losses under delegation. We consider two extreme cases. Assuming that the algorithm has no incentives under full delegation to distort her recommendations, and consequently her decisions, we expect $p = x$ and thus

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

Alternatively, we assume that the algorithm does not fully re-optimize and but continues to misrepresent her information. In this worst-case scenario, we have $p = x/\lambda$ and

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

Table 7 summarizes our counterfactual estimation results. For all hotels, we estimate that the status quo can mitigate only 1 – 4% of the profit loss from complete inaction in a volatile environment. These gains are even lower if the adjustment and information acquisition costs a manager incurs for achieving these gains are taken into account. This finding is in line with our observations that managers update prices only rarely, with delay, and are very likely to just copy the recommendation.

Overall, delegating pricing to the algorithm is likely to improve outcomes significantly. We estimate that a hotel which fully delegates to an unbiased algorithm would see a reduction of 8 to 50 percentage points in the losses accrued from mispricing as shown in Table 7. Roughly 80% of this improvement comes from more frequent price adjustments, while 10% depends on the algorithm reporting truthfully. The potential gains from delegation are however not the same for each hotel. For hotels A, B and I, for example, delegation would leave significant surplus on the table because we estimate that most of the variation in optimal pricing can only be discovered by the local hotel manager.

8 Conclusion

Algorithmic recommendations are used extensively to support decision making in organizations. In this paper, we provide a framework to understand the strategic interactions between automatic recommendations and human decisions. The crucial friction in our

model originates in managerial inattention, leading to biased communication and 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. This insight is likely relevant in a variety of other economic settings and especially important in environments in which the decision maker has a status quo option that does not require active participation. Examples of such settings include recommendations systems for parole decision (Berk, 2017), monitoring adherence to government regulation (Glaeser et al., 2021) and restocking inventory (Shang et al., 2008).

As Ludwig and Mullainathan (2021) argue, 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. Strategically responding 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 recommendation is attempting to debias her decisions.

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. We believe the answer has to do with the perception that algorithmic systems are biased, as in

our case, and that they are likely to make costly mistakes. [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

- 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,” *The Quarterly Journal of Economics*, 2011, 126 (4), 1909–1960.
- Assad, Stephanie, Robert Clark, Daniel Ershov, and Lei Xu**, “Algorithmic pricing and competition: Empirical evidence from the German retail gasoline market,” 2020.
- 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.
- Brown, Zach Y. and Alexander MacKay**, “Competition in pricing algorithms,” *American Economic Journal: Microeconomics*, forthcoming.
- 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.
- Chan, David C., Matthew Gentzkow, and Chuan Yu**, “Selection with variation in diagnostic skill: Evidence from radiologists,” *The Quarterly Journal of Economics*, 2022.
- 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.
- 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.
- Garcia, Daniel, Juha Tolvanen, and Alexander K. Wagner**, “Dynamic demand estimation using managerial responses to automated price recommendations,” *Management Science*, 2022.
- Glaeser, Edward, Andrew Hillis, Hyunjin Kim, Scott Duke Kominers, and Michael Luca**, “Decision authority and the returns to algorithms,” 2021. Harvard Business School Working Paper.
- Hoffman, Mitchell, Lisa B. Kahn, and Danielle Li**, “Discretion in hiring,” *The Quarterly Journal of Economics*, 2018, 133 (2), 765–800.
- Huang, Yufeng**, “Pricing Frictions and Platform Remedies: The Case of Airbnb,” *Available at SSRN 3767103*, 2022.
- 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,” *The Quarterly Journal of Economics*, 2018, 133 (1), 237–293.
- 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*, forthcoming.
- Nakamura, Emi and Jón Steinsson**, “Five facts about prices: A reevaluation of menu cost models,” *The Quarterly Journal of Economics*, 2008, *123* (4), 1415–1464.
- Ribers, Michael A. and Hannes Ullrich**, “Battling antibiotic resistance: Can machine learning improve prescribing?,” 2019. 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.
- 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.

Table 6: Parameter Estimates of Model

Hotel	σ_x	σ_y	σ_z	σ_c	b_1	b_2
Pooled	0.038 (0.001)	0.053 (0.001)	0.019 (0.003)	2.28 (0.03)	0.090 (0.015)	0.897 (0.029)
A	0.038 (0.001)	0.100 (0.004)	0.007 (0.009)	1.34 (0.03)	0.038 (0.010)	0.354 (0.021)
B	0.023 (0.001)	0.056 (0.002)	0.024 (0.008)	2.06 (0.05)	0.032 (0.011)	0.644 (0.028)
C	0.026 (0.001)	0.034 (0.001)	0.026 (0.001)	2.25 (0.02)	0.039 (0.006)	0.383 (0.017)
D	0.017 (0.001)	0.037 (0.001)	0.014 (0.003)	2.64 (0.06)	0.035 (0.012)	0.676 (0.027)
E	0.028 (0.001)	0.032 (0.001)	0.040 (0.002)	1.79 (0.02)	0.016 (0.002)	0.273 (0.009)
F	0.019 (0.001)	0.036 (0.001)	0.001 (0.005)	1.84 (0.04)	0.027 (0.007)	0.397 (0.025)
G	0.032 (0.001)	0.026 (0.001)	0.035 (0.002)	2.54 (0.03)	0.078 (0.008)	0.809 (0.025)
H	0.035 (0.002)	0.067 (0.004)	0.063 (0.009)	1.47 (0.03)	0.007 (0.005)	0.516 (0.015)
I	0.034 (0.002)	0.083 (0.003)	0.08 (0.005)	1.54 (0.03)	0.034 (0.011)	0.644 (0.024)

Notes: Estimated parameter values for each hotel, A to I, and pooled across hotels. Bootstrapped standard errors in parenthesis.




Figure 6: Bias Parameter Validation

Notes: Each point represents a hotel. The x-axis gives the bias parameter identified using the optimality condition for the algorithm and the y-axis gives the bias parameter using only information of the hotel manager.

Table 7: Counterfactuals

Hotel	Benchmark	Profit Loss	Delegation	Biased
A	0.991	0.973	0.874	0.875
B	0.990	0.977	0.875	0.880
C	0.984	0.967	0.727	0.734
D	0.984	0.970	0.835	0.851
E	0.986	0.970	0.761	0.779
F	0.996	0.990	0.774	0.800
G	0.984	0.968	0.643	0.673
H	0.990	0.976	0.872	0.891
I	0.994	0.983	0.920	0.921

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 last 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.

Supplementary Materials for Online Publication

for

“Strategic Responses to Algorithmic Recommendations: Evidence from Hotel Pricing”

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

A Proofs

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,

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

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 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 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

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

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

Proof of Proposition 1. Notice first that

$$\begin{aligned} (r - \tilde{x}(r) - y)^2 + \sigma_z^2 &< b_2c \\ \Leftrightarrow x - r - \sqrt{b_2c - \sigma_z^2} &< y < x - r + \sqrt{b_2c - \sigma_z^2}. \end{aligned}$$

Denote $d(c) := \sqrt{\min\{b_2c - \sigma_z^2, 0\}}$. Then we can write

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

If $b_2\sigma_y^2 < b_1\sigma_z^2$, then $b_2c(0) < \sigma_z^2$, $d(c(0)) = 0$, 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 $r \rightarrow \infty$, the integrand vanishes, while the denominator converges to 1. Hence, $\lim_{r \rightarrow \infty} \mu(r) = 0$. Finally, to see that $\mu(r)$ is increasing in a neighborhood of $r = 0$ observe that

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

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

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

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. \square

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

$$\begin{aligned}E(y \mid r, y \notin Y_0(r, c)) &= \frac{1}{A} \int_0^{c(r)} \int_{y \notin Y_0(c, r)} y d\Psi_y(y) dF(c) \\ &= \frac{1}{A} \int_0^{c(r)} \left(\int_{-\infty}^{r-\tilde{x}(r)-d(c)} y d\Psi_y(y) + \int_{r-\tilde{x}(r)+d(c)}^{\infty} y d\Psi_y(y) \right) dF(c),\end{aligned}\quad (4)$$

where

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

Notice then that,

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

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. \square

Proof of Corollary 1. It is enough to show that

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

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. \square

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

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

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

Rewriting we have

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

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,

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

\square

B 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

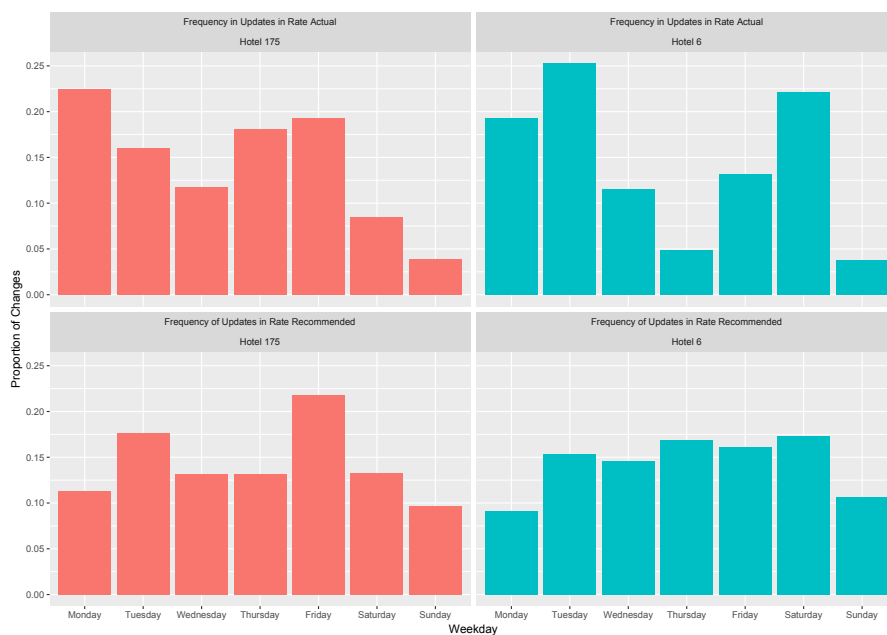


Figure 7: Frequency of Updates in Prices and Recommended Rates Across Weekdays for Hotel 6 and 175. Source: [Garcia et al. \(2022\)](#).

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 7 plots the relative frequency of price changes and recommendation changes for the two biggest hotels in the sample. The main takeaway from this picture is that the hotel 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 8 which shows what share of all price updates are done on each weekday in each of the hotels. 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 10 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 when the manager copies the recommendation for at least a full arrival week for at least

Table 8: 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. \(2022\)](#).

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 and discussion about the adjustment costs is provided in [Garcia et al. \(2022\)](#).

C 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 r , $Corr(p, r \mid p \neq r) = \lambda$ and

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

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

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

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) and 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 scalars of the distribution.

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

D 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 manager’s signal realizations for 10 evenly split percentiles starting from the 50th percentile. The results are reported in Table 9. 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 percentage of the best response outcome. Since, high signal realizations are the ones where the algorithm already has a very high chance of inducing the hotel to revise its price even without exaggeration, we also estimate the results for the 99th percentile.

Table 9: Percentage changes in actions and losses relative to best response

	Percentile of F_x										
	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.85	0.9	0.95	0.99
Δr	100%	1.0%	3.3%	-15.4%	2.5%	-2.6%	-4.9%	-0.8%	-14.0%	-16.4%	-19.1%
Δloss	-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 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 managers 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 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. $100\% \times \frac{\mathbb{E}[(p(r^*(x)) - p^*)^2] - \mathbb{E}[(p(x/\lambda) - p^*)^2]}{\mathbb{E}[(p(r^*(x)) - p^*)^2]}$.

For most signal realizations the difference in the advisor’s payoff when best responding compared to when playing the linear strategy is negligible. For all but the two highest percentiles 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 generate significant incentives for deviating and hence is likely to have a negligible quantitative impact on the results.

E Alternative Models

We define a recommendation $r(x)$ to be unbiased if

$$E[p^* | x] = r.$$

We now discuss a handful of alternative models to the one presented in the main text.

- Suppose that p^* is normally distributed with mean zero and standard deviation σ . Furthermore, assume that the manager observes signal y and the 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 recommender 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

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

In our data $E[p^* | p \neq r] = \gamma r$ 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,

$E[p^* | x, x \neq p^*] = \gamma x$. In this case, we have that the recommendation is unbiased only if $r = (q + (1 - q)\gamma)x$ in which case

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

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 8, but only 13% of these updates result in $r' = p$ and $E(r' | p, p < r) > p$, i.e. the recommendation does not fully react to the change in price as shown in Figure 9. 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 (see Figure 10) while not changing a single price manually.
- **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.




Figure 8: Probability of the recommendation changing after the price manually changed to not equal the recommendation. Ticks represent the 95% confidence intervals.

Huang (2022) shows that Airbnb’s recommendations to hosts are downward biased, consistent with their preference for revenue-maximization. By contrast, we do not observe a significant directional bias in our setting: the algorithm exaggerates both price hikes and drops by approximately the same amount.




Figure 9: 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 10: Number of Rate Updates (in Log) Conditional on Copying and Not Copying Recommended Rates. Source: [Garcia et al. \(2022\)](#).