CESifo
AREA CONFERENCES 2021

Economics of Digitization
Munich, 19 – 20 November 2021

Algorithmic Recommendations and Human Decisions: Evidence from Hotel Pricing

Daniel Garcia, Juha Tolvanen, Alexander K. Wagner

CESifo

Algorithmic Recommendations and Human Decisions: Evidence from Hotel Pricing*

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

October 18, 2021

Abstract

We study the interaction between algorithmic advice and human decisions in the context of dynamic pricing in hotels. We provide evidence of adjustment costs and inertia that induce a conflict of interest between the human manager and the algorithmic advisor, resulting in biased recommendations and suboptimal decisions. We develop a simple structural model of advice that is consistent with the data and allows us to quantify the losses from this bias as well as the benefits from a potential shift into a fully automated algorithmic pricing system.

Keywords: D11, D83, L13

JEL-classification: Advice; Algorithmic Recommendations; Delegation.

---

*

⊕ 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

⊗ Vienna Center for Experimental Economics and Department of Economics, University of Vienna, Oskar-Morgenstern-Platz 1, A-1090 Vienna, Austria. Email: alexander.k.wagner@univie.ac.at

1 Introduction

Organizations rely increasingly on algorithmic systems to make decisions. Examples range from hiring policies at tech firms (Hoffman et al., 2018) to bail decisions by judges in Florida (Kleinberg et al., 2018). Algorithmic systems are particularly prominent in pricing, both in online platforms (Chen et al., 2016) as well as in traditional markets (Assad et al., 2020). In a large fraction of these applications, however, human managers retain final decision rights and machines act as mere advisors. As the literature on advice shows,[1] algorithmic advice will be effective only to the extent that the incentives of the human and the machine are sufficiently aligned. In this paper, we study algorithmic recommendations in hotel pricing, identify a source of misalignment, and show that it leads to suboptimal decisions.

To do so, we leverage a proprietary dataset that contains approximately 5.5M prices for a number of hotels over a period of just under a year. Hoteliers have access to algorithmic price recommendations provided by an external revenue management services company. Because of time constraints, managers change prices only infrequently (roughly once a month). Algorithmic recommendations change much more often, which inevitably leads to divergence between prices and recommendations. We show that this wedge induces a moral hazard problem since the algorithm would like to motivate the manager to change prices more frequently. We show that the algorithm optimally *exaggerates* its private information to incentivize faster updating, leading to substantial welfare losses. We finally perform a counterfactual exercise where decisions are fully delegated to the machine and show that it significantly outperforms the current arrangement for most hotels.

We begin our analysis by establishing three empirical patterns. First, price updates are rare. On average, managers update prices for a particular room-arrival-date combination once every month. Second, managers often update the price to exactly match the recommended price (85%), and they are even more likely to do so when the recommendation's cumulative change since the last update is large (95% if it exceeds 10%). Third, whenever

[1]See Sobel (2013) for an excellent summary of the vast literature. In Section 2 we discuss the contributions of our study to this literature.

managers depart from the price recommendation they still tend to incorporate it in the price, but they do so imperfectly. A one-percent increase in the recommendation leads to a 0.72 percent increase in the expected price, provided that the two are different.

To make sense of these findings, we build a simple model of information processing that mimics the environment faced by our managers. In our model, managers display limited attention to price changes and face adjustment and information processing costs.[2] We assume that the manager observes (for free) the current price and the price recommendation for that particular product (a specific room-arrival-date combination). She then decides whether to keep the current price (at no additional cost) or to put some attention to this specific price (at a positive cost). If she chooses the latter, she gets to observe a private signal about the specific demand and cost conditions of this product. At that point, the manager has to decide whether to adjust the price to equal the recommendation (at no additional cost) or to proceed further and acquire more information before making her final decision (incurring a further cost). This timeline is motivated by the distinctive features of the dashboard displayed to the hoteliers, which makes it easy to implement the recommendation, and requires hoteliers to access a different tab in order to manually update the price.

The model is able to match the empirical patterns discussed above. For each realization of the recommendation, there exists a cutoff cost such that the manager devotes attention to this price only if her cost shock is lower than that cutoff. This cutoff is increasing in the size of the recommendation, as larger deviations motivate higher effort. As a result the expected cost draw, given that the manager puts attention to a price is increasing in the recommendation size. This induces positive selection and helps explain the increased likelihood of matching the recommendation as the cumulative change grows. Finally, if the hotelier believes that the recommendation is biased (exaggerated), she will be more

[2]There is a large literature in macroeconomics that studies adjustment and information processing costs in price setting. Our model is closest to Alvarez et al. (2011), which includes both types of frictions. The main difference is that we study a static model. As we argue in Section 4, the data is consistent with static optimization.

likely to depart from it when her private information suggests a change in a different direction (say increase rather than decrease). Thus the hotel manager is more likely to set her own price when her private information contradicts the direction of change in the recommendation. This pattern, which we call negative selection, further dampens the pass through. To sum up, the model is able to generate the somewhat counterintuitive pattern of a high unconditional copy rate combined with considerably dampened pass-through of recommendations when the hotelier sets her own price.

The crucial component of the model is, therefore, the perceived bias in the recommendation. We rationalize this bias as the endogenous outcome of the interaction between the algorithm (or its coder) and the hotelier. We posit that the algorithm wants to induce optimal decisions and has correct expectations about the hotelier's response to different recommendations. The algorithm secretly chooses a reporting strategy which, for simplicity, we assume to be a linear factor that multiplies the change in the privately observed component of the price. If this factor is above one, the algorithm would exaggerate her private information, while if it was below one it would be dampening it. In equilibrium, the hotelier has correct beliefs about this factor and, therefore, forms correct expectations about the private information held by the algorithm. Because the hotelier often chooses a price that exactly matches the recommendation, this distortion translates into suboptimal decision-making.

The model provides a good fit for the data. We recover the parameters using a minimum distance estimator and imposing the equilibrium condition on the resulting bias component. For the whole sample of hotels, the bias factor we identify is 1.25, which is significantly higher than one but also lower than the naive estimate of 1.38 necessary to explain the low pass-through in a model without selection. There is, however, considerable heterogeneity across hotels with estimates ranging from 1 to 1.5. We also find that hoteliers' information is more volatile than the information held by the algorithm, but since this information is only rarely used it shows very high costs of information processing. This translates into inaccurate decisions: the standard deviation of the difference between the optimal price and the realized price is larger than the standard deviation of the price we observe in the data.

A natural implication of these results is that delegation may improve decision-making. Delegation has some obvious benefits (instantaneous and costless decisions) but also some costs: the information privately held by the manager is no longer part of the decision. In addition, it eliminates the moral hazard friction and therefore results in aligned preferences. We show that for a majority of hotels in our sample, delegation would reduce losses by about 10%. Out of these gains, the majority comes from more frequent decisions (70%), while the rest comes from de-biasing (15%) and costless information processing (15%).

2 Related Literature

There exists a vast theoretical literature on strategic advice in economic organizations.[3] 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 in that, conditional on the realized state of the world, they disagree on the 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 decision makers' span-of-control may greatly exceed that of the agents with whom they communicate. In the context of the interaction between human decision-makers and machine advisors, this is almost certainly guaranteed to occur, as the attention cost of the machine is infinitesimal compared with that of the human. In the spirit of Aghion and Tirole (1997) machines hold real authority as the principal's information processing costs greatly exceed those of the algorithm.

A similar tension arises in (Kartik et al., 2007), who introduce a cheap talk model in which a fraction of the audience is *naive* and takes the message at face value. In

[3]See Sobel (2013) for a recent survey on communication, focusing mostly on cheap-talk messages. Kamenica (2019) provides 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.

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 case, all receivers are sophisticated, but behave as if they were naive in order to economize attention costs. Our structural model includes this trade-off but takes into account that the attention decision is endogenous.

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 a setting where the 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 will tend to focus her attention on those with higher stakes and fully delegate decision-making to the machine in those with lower stakes. In our setting, we observe hoteliers choosing a price that exactly matches the recommendation (akin to delegation, or rubber-stamping) more often precisely when the recommendation departs *less* from the status quo. The key insight, however, is that even rubber-stamping requires some attention from the principal as opposed to fully automated decision-making.[4]

Bloedel and Segal (2020) study a persuasion model where the principal is subject to rational inattention. Just like in our setting, inattention induces a sort of moral hazard problem that leads the advisor to distort her messages as a way to motivate the principal to pay attention. They show that it is optimal to fully disclose messages only when 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.[5]

More broadly, we contribute to the empirical literature on strategic communication

[4]In the algorithmic pricing industry, some companies offer an arrangement similar to that suggested in Agrawal et al. (2019). The algorithm directly implements changes if they are within a certain range of the target price, while human approval is needed if the suggested price falls outside that range.

[5]We do not observe any form of pooling or bunching in the data, whether in levels or in logs. On average, prices set manually by the hotelier increase continuously in the recommendation.

in organizations.[6] We are aware of two papers that use equilibrium analysis to identify strategic communication behavior. Backus et al. (2019) provides evidence of strategic communication in bargaining in an online marketplace, where 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 a number of advantages. First, both the sender and the receiver are professional and face serious financial consequences from their actions. Second, there is an obvious mapping between messages and recommended actions. Third, the action space is sufficiently rich that we can directly identify the posterior beliefs of the receiver whenever she chooses a price that departs from the recommendation.

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, and assess potential advantages (accuracy, speed) and disadvantages (algorithmic bias or negative perception of third-parties).[7] 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. Just like in our setup, human inertia is a major concern but the algorithmic recommendation is assumed to be non-strategic. Our study suggests that assuming truthful recommendations as a counterfactual scenario may be neither optimal nor realistic.

3 Data and Institutional Background

Our data set contains more than 5 million observations of hotel-room pricing information and the corresponding universe of about 60 thousand bookings, all aggregated at the

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

[7] Some examples of papers in this vein include Hoffman et al. (2018), Kleinberg et al. (2018), Ribers and Ullrich (2019), Chan Jr et al. (2019) and Currie and MacLeod (2017).

daily level. This high-resolution proprietary data set was made available by our corporate sponsor, whose identity is withheld by request. The corporate sponsor is based in Central Europe and provides a number of revenue management services, including price recommendations, to a large number of independent hotels. The rate and booking data comes from 9 different hotels, eight of which are located at resort destinations (A-H), and one in an urban area (hotel I). Our data contains bookings and prices for each hotel over a period of about 14 months. For each potential arrival date, we observe the flow of bookings, actual rates charged by the hotel, and rates recommended for every room by the revenue management company. The actual rate can be thought of as an index price set by the hotel. It is modified by pre-specified and channel-specific discounts or fees and is then pushed to each channel in which the product is offered. Both the revenue management system (which we shall refer to as the algorithm) and the hotelier use this variable as the main instrument for price optimization.[8]

A crucial input to our analysis is the recommended rate provided by the algorithm. The revenue management company offers an algorithmic price recommendation service with the objective to maximize the client's revenue through optimized pricing. For price optimization, the revenue management company has access to all bookings from a hotel on a real-time basis, as well as additional demand-related information such as local variation in weather conditions, events, public holidays, hotel reputation, competitor prices, etc. In addition, the revenue manager and the hoteliers are regularly exchanging information about local demand conditions. Hoteliers are trained to transmit private information about local demand shocks to the revenue management company. The company combines all relevant information and feeds it into its proprietary pricing algorithm to produce a rate recommendation for each product of the hotel.[9]

[8] See Garcia et al. (2021) for more details.

[9] We do not have access to the company's proprietary pricing algorithm. A common concern in competition policy is that algorithmic pricing can be used to coordinate with competitors. We believe that coordination through algorithmic pricing is not much of a concern in our context. Given the highly fragmented nature of the industry, the number of hotels that use recommendations from our corporate sponsor is very small ($\leq 1\%$) compared to the size of the industry in the considered regions.

The hotelier decides every day whether to log into the system. If she does, she observes the current rate recommendation for each room and then decides which rates to update and by how much. After manual confirmation by the hotelier, updated actual rates are pushed from the algorithm to the hotel's property management system and distribution channels (mostly OTAs). Although the hotels in our data are typical for their respective regions, with about 50 rooms per hotel they are relatively small by international standards. Most of them are family-run which means that the manager's job description includes many responsibilities beyond revenue management. From private communications with the revenue management company, we learned that managing prices takes only a small fraction of a hotelier's weekly work schedule.

For most of our analysis, we will rely on recommendation and price changes as the main variables (see Table 1). From the panel of daily rates, 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 recommendation to be the change in the logarithm of the algorithm's recommended price since the last update. We restrict the sample to include only those observations in which the initial price matched the recommendation. This allows us to interpret differences between the final price and the recommendation as differences in the current information processing. It has the additional benefit of removing any feedback effects from past prices into future recommendations. The resulting sample includes approximately 52% of observations and 64% of the price updates.

4 Stylized Facts and Relationship to Literature

In this section we present a selection of empirical patterns from the data that we think are important and that allow us to pin down the right model. In this section we pool together observations from all hotels but these features hold (qualitatively) for all hotels in our sample.

Prices adjust infrequently relative to recommendations

The first observation is that actual rates change only infrequently. On average, we observe a price change every 35 days, but there is considerable heterogeneity across different hotels (see Table 1). Algorithmic recommendations change much more frequently (every week or so), which inevitably leads to a disparity between both. This inertia is also reflected in the distribution of price changes, as there are relatively few observations around 0 (see Figure 2). Both of these patterns indicate that the hoteliers are facing considerable adjustment and information costs (Nakamura and Steinsson, 2008).

The update rate is positively correlated with the change in recommendation

Larger changes in the recommendation are associated with a higher likelihood of a price update, as depicted in Figure **??**. For instance, if the algorithmic recommendation has not changed since the last time that the hotelier updated the price, the probability of an update is less than 1%, while if the current recommendation is not within a ten-percent band of the original, this probability exceeds 11%. This implies that recommendations contain important information. Furthermore, the adjustment costs that the hoteliers face are not limited to opportunity cost of sitting down at the computer to check the recommendation. The hotelier must face an additional adjustment cost after observing the recommendation to choose not to update the price every time they observe the recommendation changing.

The probability that the updated price matches the recommendation is increasing in the change in the recommended price

Conditional on an update, hotel managers are very likely to update the price to exactly match the recommendation. On average, around 85% of the price changes result in a price that equals the current recommendation, with small differences across hotels. This percentage is even higher if we condition on a large change in the recommendation (see Figure 4). In particular, whenever the recommendation change exceeds 10%, the hotelier will choose a price that exactly matches the recommendation with 95% probability. To the extent that the algorithm would prefer the manager to change prices more often and more intensively (see below), this pattern is inconsistent with off-the-shelf models of advice, whereby the influence of the adviser decreases when making more *extreme* recommendations .[10]

The hazard rate of updates is (approximately) constant

There is a large literature on price adjustment in macroeconomics, whereby firms engage in forward-looking optimization. Most of the literature fits into two camps: exogenous and endogenous adjustment decisions. If the timing of adjustment is exogenously determined we expect the time since the last update to follow a exponential distribution (as discussed in Alvarez et al. (2016)). If instead the adjustment is endogenous, we would typically observe an initial inaction period, followed by a discrete jump in the probability. In our data, the hazard rate of updates is almost perfectly constant, aligning with the predictions in the first camp (see Figure 3). Because the cross-sectional variation of updates can only be explained by endogenous adjustment (see above), we will consider a static optimization framework.[11]

The changes in the recommendation are only partially passed through to realized prices

If the interests of the hotelier and the revenue manager are perfectly aligned and the hotelier's arrival of private information is uncorrelated with the direction of that private information, one would expect that on average a one Euro increase (decrease) in the recommendation would bring about a one Euro increase (decrease, respectively) in the price. The realization of the difference between the two would then represent the additional, idiosyncratic information held by the manager.[12] Instead. we observe a pass-through rate of around 72% (see Table 2). In other words, when hoteliers manually update their prices the

[10]This 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). strategic concealment and disclosure (Aghion and Tirole, 1997) this leads the principal to maintain decision rights.

[11]Notice that the typical hotel in our data has roughly 2000 different products (different rooms for different dates) on a given day. It is probably very difficult for her to keep track of her previous updates and framing the problem as static is likely to provide a good approximation.

[12]See Appendix ? for details.

average price change dampens the recommended price change. Including various controls (room type-arrival week fixed effects and a polynomial of the days before arrival) leads to a modest increase in the estimated coefficient (73%). It follows that hoteliers' must believe that the revenue managers' algorithm exaggerates the optimal price change (on average).

Interestingly, the (unconditional) relation between recommnedations and prices is almost perfectly linear (see Figure 1). This fact is inconsistent with equilibria in traditional cheap-talk models, which display discontinuities/bunching to ensure incentive compatibility. It is also at odds with mutlidimensional 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 actually copies it.

5 Model

We now introduce a simple model of price adjustment with recommendations and private information acquisition that can rationalize the findings above. 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. Hoteliers care about profits, defined as $\Pi = \Pi_0 - \eta(p - p^*)^2$, where p^* is the ideal price given demand and cost conditions and $\eta > 0$ is a parameter.[13] We assume that $p^* = x + y + z$, with x, y, z being symmetric, mean-zero random variables, with respective variances σ_i^2, for $i = x, y, z$. In our empirical

[13]This expression can be micro-founded as the profit function with a log-linear demand, with semi-elasticity η.

specification we will assume that these are normally distributed, but none of the theoretical results depend on this feature. We assume that all these variables are drawn independently. Variable x is privately observed by the algorithm (but not by the hotelier). The hotelier does observe some recommendation r, which we assume is a linear function of x: $r = \frac{1}{\lambda}x$, for some $\lambda > 0$. Variables y and z may be observed by the hotelier (upon incurring their corresponding costs costs).

We assume that the manager observes the current price (normalized to 0) and the recommendation r, as well as the realized adjustment costs (c_1, c_2). Upon observing this information, the hotelier may decide whether to stop and maintain the current price ($p = 0$) without incurring any additional cost or acquire signal y at a cost c_1. We think of y as information that the hotelier may gather from her own recollections of events concerning a particular product (e.g. the cook is on vacation that week) and does not require further information acquisition efforts (e.g. accessing booking histories, competitor prices, etc.). If the hotelier does observe y, then she faces the choice whether to update the price to exactly match the recommendation, which she can do directly in the dashboard, resulting in $p = r$ or accessing further information z at cost c_2. Only at this later stage can the manager update the price freely ($p = E(p^* \mid r, y, z) = E(x \mid r) + y + z$).

In order to keep the model as parsimonious as possible, we assume that costs c_1 and c_2 are both directly determined by the realization of a cost shock c, drawn from a distribution $F(c)$. In particular, we assume that $c_i = b_i c$, with $b_i > 0$. This cost structure is consistent with two different interpretations. One may view c as the opportunity cost of a unit of time for the hotelier and b_i is then the time spent in each of the two stages. Alternatively, c is a cognitive cost shock while b_i measures the complexity of the task. We assume that c is observed by the hotelier before entering the decision-making process.

The crucial ingredient of the model is the recommendation. We assume that the algorithm chooses λ before observing x in order to maximize profits (Π). Crucially, the algorithm does not care about the costs incurred by the hotelier when updating the price, introducing a conflict of interest. The hotelier does not directly observe λ but has some belief, which we impose will be correct in equilibrium. Let $\tilde{x}(r)$ denote the induced expectation, given some belief about λ.

5.2 Analysis

We begin our analysis with the problem of the hotelier. Upon observing (r, c) the hotelier decides whether to initiate the information acquisiton 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 r.

Alternatively, she may continue which has some expected loss $l(r, c)$. In order to characterize this expected loss we need to solve the problem once the hotelier acquires information y. In such a case, the hotelier will choose to match the recommendation whenever $(r - \tilde{x}(r) - y)^2 + \sigma_z^2 < c_2$. Let $Y_0(r, c)$ denote the set of values of y for which this expression holds. Notice that $Y_0(r, c)$ is an interval centered at $r - \tilde{x}(r)$. The expected loss is then:

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

where Ψ_y is the cumulative distribution function of a mean-zero normal distribution with variance equal to σ_y^2.

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,*

$$l_1(r, c) \leq 2r(1 - \tilde{\lambda})^2 \int_{y \in Y_0(r,c)} d\Psi_y(y) \text{ for all } r > 0.$$

If $\sigma_y > 0$ and $\tilde{\lambda} < 1$, the loss functions is strictly increasing in $|r|$ and both of the above 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 hotelier will continue its information acquisition effort for costs lower than $c(r)$ and will keep the old price for costs higher than $c(r)$. Notice that $c(r)$ is an even function. Conversely, for a given cost, the set of recommendations such that an update occurs can be written as the union of two intervals $(-\infty, -\bar{r}(c))$ and $(\bar{r}(c), \infty)$, for some function $\bar{r}(c)$. The following lemma provides a condition such that that both $\bar{r}(c)$ and $c(r)$ are strictly increasing functions.

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

In other words, higher cost realizations require larger deviations in the recommendation before the hotelier will re-evaluate 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 observations described above. Importantly, however, the probability of an update itself is increasing in $|r|$. This implies that bigger changes in the recommendation are associated with a higher chance that the hotelier considers copying the recommendation in the first place. The resulting relationship between the size of the recommendation change and 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.

Proposition 1. *If* $b_2\sigma_y^2 < b_1\sigma_z^2$, *then* $\mu(r) = 0$ *for all* r. *Else,* $\mu(0) > 0$, $\lim_{r\to 0} \mu(r) = 0$, *and* $\mu'(r) \geq 0$ *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 hoteliers 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. Second, if $F(c(0)) = 1$, there is no inertia (hoteliers update the price every period) and higher r makes copying the recommendation less appealing (since it is associated with a higher loss). This comparative static holds in models of strategic delegation, whereby the principal (hotelier) is more likely to rubber-stamp *low* recommendations (e.g. Aghion and Tirole).

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

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

Since, $Y_0(r,c)$ is centered at $r - \tilde{x}(r)$, $E(y \mid r, y \notin Y_0(r,c))$ depends on $\tilde{\lambda}$. If $\tilde{\lambda} = 1$, then $r - \tilde{x} = 0$, $Y_0(r,c)$ is centered at the origin and hence $E(p \mid r, y \notin Y_0(r,c)) = \tilde{x}(r) = r$. Instead, if $\tilde{\lambda} \in (0.5, 1)$, then 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 $\tilde{\lambda}$. This means that we cannot directly identify $\tilde{\lambda}$ from the pass-through rate but instead must be jointly identified with the remaining parameters.

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

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

and

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

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

Corollary 1. *Conditional on the hotelier not copying the recommendation but changing the price, the hotelier's private information is negatively correlated with the recommendation, i.e.*

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

We finally discuss the problem of the algorithm. The algorithm chooses λ to maximize expected profits but this choice is unobserved by the hotelier. 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 hotelier to change prices less frequently. Second, it has an ambiguous impact on the probability that the hotelier chooses a price that exactly matches the recommendation, as the function $\mu(r)$ is non-monotone. Third, it reduces the distance between the recommendation and the optimal price which translates into increased profits.

6 Estimation and Results

For the empirical implementation, we assume that c 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 (σ_c, b_1, b_2). Additionally, we need to estimate the reduced-form parameter $\tilde{\lambda}$ which measures the equilibrium bias in the revenue manager's recommendations. To estimate these parameters, we use a method of simulated moments, minimum distance estimator with seven target moments (see Table 3) that additionally imposes the restriction that there is no (secret) profitable deviation from the recommendation algorithm for the revenue manager. Four of those targets (1 to 4) 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.

In practice our estimation algorithm proceeds as follows. We first fix a level of bias $\tilde{\lambda}$ and simulate the model to find structural parameters 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 manager's payoff. If such a profitable deviation exists, we pick such a deviation $\tilde{\lambda}'$ 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 and in 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. We estimate the model both for the pooled data and for each hotel individually.

We now discuss identification. To this end it is convenient to consider a special case of the model in which $\sigma_y = 0$. 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 (given λ). Likewise, the ratio of the copy rate for large recommendation changes over the average rate determines the standard deviation of the cost distribution. Given this, 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 moments are closely linked to each of those parameters, with the standard deviation of price changes now helping to determine the variance of y.

For a given λ and $F(c_i)$, σ_x is identified by the observed distribution of recommendations ($Var(r \mid U)$). Given σ_y, and λ, σ_z is directly pinned down by the standard deviation of the difference between the price and the recommendations ($Var(p - r \mid U)$). Now fix the joint distribution of signals and notice that b_1, b_2, σ_c are then pinned down by the three last moments. The last two moments then determine λ and y jointly.

We run the routine on the pooled dataset first. Results are summarized in the first row of Table 4, including bootstrapped standard errors. We find that the private information held by hoteliers accounts for less than 20% in the total variance of the optimal price. That is, managers' private information is at least five times as valuable as that of the algorithm. Unfortunately, accessing this information requires substantial effort. We estimate a mean adjustment cost (c_1) of approximately 0.2, with a standard deviation of 1.10.[14] The cost of acquiring further information is estimated to be an order of magnitude larger (mean 3.06 and s.d. 4.24). As a result, there is considerable dispersion between realized and (counterfactual) optimal prices.

Our estimates also suggest a significant bias in recommendations ($\lambda = 0.805$). Since most of the updates match the recommendation exactly, this translates into biased prices.

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

Nevertheless, the welfare impact of this bias is ameliorated by the fact that the manager selects into the decision to match the recommendation (see Proposition 2).

The model is able to fit the target moments well (see Table 3). It also does a reasonably good job at replicating the empirical patterns described in Section 4. For instance, it predicts an update rate of about 15% when the recommendation exceeds 5%, which is a bit higher but reasonably close to the data. It also generates a relation between recommendations and prices (conditional on an update) that it is consistent with the data.[15]

On average, we find that the current regime does a poor job at exploiting the available information and is able to reduce losses from mispricing by less than 4%. This is not very surprising, since managers only rarely update prices and, when they do, they tend to match the recommendation.

We then run the routine separately for each hotel. We summarize results in rows 2-10 in Table 4. Most hotels are well-represented by the pooled data. The variance of x accounts for less than 30% of the total variance for all of them. There is, however, considerable heterogeneity in the contribution of y versus z and the cost distributions.

We finally compute the expected profit loss for each hotel in the status quo, relative to the profit loss they would experiment if they never updated their prices. This metric is independent of η and takes into consideration that some hotels experience a more volatile environment than others. Formally,

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

7 Counterfactual Exercise: Delegation

In light of these results, it appears that a majority of hotels would benefit from delegating decisions to the algorithm. Delegation to an algorithmic system completely eliminates

[15]Notice that the scatter plot of the model has a smaller support than the model because the kurtosis of the distribution of recommendation changes exceeds that of the normal distribution.

information processing and menu costs, thereby eliminating delay in decision-making. It should also lead to truthful recommendations since we identified inertia as the source of conflict between the algorithm and the hotelier. In exchange, the algorithm does not have access to the realizations of signals y and z, which the data suggests are even more valuable than x. The key insight, however, is that the hotelier's inertia reveals that this information comes at significant costs, greatly reducing its value for decisions.

Our parameter estimates directly allows us to compute the residual losses under delegation. We consider two extreme cases. Under the assumption that in the delegation regime the algorithm has no incentives to distort its recommendations/decisions, we would expect $p = x$, and thus, $w_i = \frac{\sigma_y^2+\sigma_z^2}{\sigma_x^2+\sigma_y^2+\sigma_z^2}$. Alternatively, the algorithm may not fully re-optimize and continue to mis-represent its information. In this (worst) case, we would see $p = x/\lambda$ and,

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

Results are summarized in Table 5. For all hotels, we estimate that the current environment can mitigate only 1-4% of the profit loss due to demand/cost volatility. This gain is even lower if we take into consideration the adjustment and information acquisition costs that the manager has to incur to achieve them. This is not that surprising since managers update prices only rarely and with delay, and they are very likely to copy the recommendation.

Delegation is likely to improve outcomes significantly. We estimate that a hotel who fully delegates to an unbiased algorithm would see a reduction of 5 to 40 percentage points in the losses due to mispricing. Roughly 80% of this improvement comes from more frequent adjustment, while 20% depends on the algorithm reporting truthfully. For hotels A and B, however, delegation would still leave significant surplus in the table since we estimate that most of the variation in optimal pricing can be discovered only by the hotelier.

8 Conclusion

Algorithmic recommendations are a prominent tool for many decision-makers. In this paper, we provide a framework to understand the interaction between recommendations and human decisions. A key friction originates in managerial inattention, leading to biased communication and decisions. We have applied our setup to a dataset containing 5M prices from hotels in Central Europe. We have shown that full delegation is likely to be welfare-improving, even if it forgoes the potential benefits of richer information.

There are many avenues for future research. The most obvious continuation of this work would involve an explicit, fully dynamic model of price adjustment.

Another question that we have not attempted to answer in the present paper is why do recommendation systems remain in place, even if they are quite inefficient. We believe the answer has to do with the perception that algorithmic systems are biased (as it is in this case), and that they are likely to make costly mistakes. Dietvorst et al. (2015) shows that when humans see algorithms err (as it happens in our setting), they tend to avoid using them, even if they are more accurate than the alternative.

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.

Alvarez, Fernando, Herve Le Bihan, and Francesco Lippi, "The real effects of monetary shocks in sticky price models: a sufficient statistic approach," *American Economic Review*, 2016, *106* (10), 2817–51.

Assad, Stephanie, Robert Clark, Daniel Ershov, and Lei Xu, "Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market," 2020.

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.

Bloedel, Alexander W and Ilya Segal, "Persuading a Rationally Inattentive Agent," Technical Report, Working Paper 2020.

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.

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

Chen, Le, Alan Mislove, and Christo Wilson, "An empirical analysis of algorithmic pricing on amazon marketplace," in "Proceedings of the 25th international conference on World Wide Web" 2016, pp. 1339–1349.

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.

DellaVigna, Stefano and Matthew Gentzkow, "Persuasion: empirical evidence," *Annu. Rev. Econ.*, 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. Forthcoming*, 2021.

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

Jr, David C Chan, Matthew Gentzkow, and Chuan Yu, "Selection with variation in diagnostic skill: Evidence from radiologists," Technical Report, National Bureau of Economic Research 2019.

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

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.

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.

Sobel, Joel, "Giving and receiving advice," *Advances in economics and econometrics*, 2013, *1*, 305–341.

9 Tables and Figures (in main text)

Table 1: Updates and Recommendations

	Min	Mean	Max
Update Rate	0.012	0.038	0.463
Update Rate (Rec 0)	0.012	0.038	0.463
Update Rate (Large Rec)	0.018	0.143	0.658
Copied Rec	0.757	0.840	0.894
Copied Rec (Large Rec)	0.869	0.950	1
Update Size	0.033	0.048	0.066
Update Size (Copy)	0.029	0.044	0.061
N	59719	223469	404196

Notes: For all statistics we report the maximum, minimum and average value across hotels. Update rate is the proportion of products in which we observe an update on a given day. We report the update rate unconditionally, conditional on the recommendation not having changed and conditional on the recommendation having changed by more than 10% (rows 1-3). The Copied Rec rate is the proportion of updatess in which we the price matches the recommendation exactly. We report both the unconditional and the value conditional on an absolute change in the recommendation of at least 10%. The update size is the average change in the price following an update. We report this value unconditionally and conditional on it matching the recommendation.

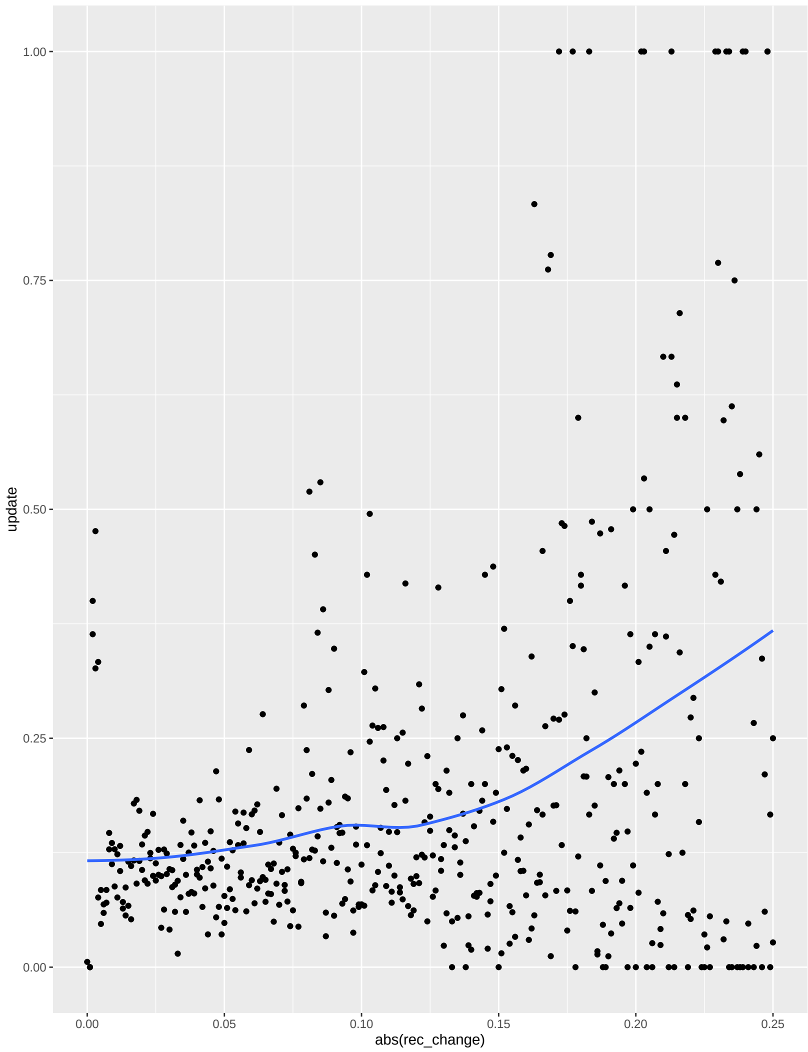

Figure 1: Update Rate

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.

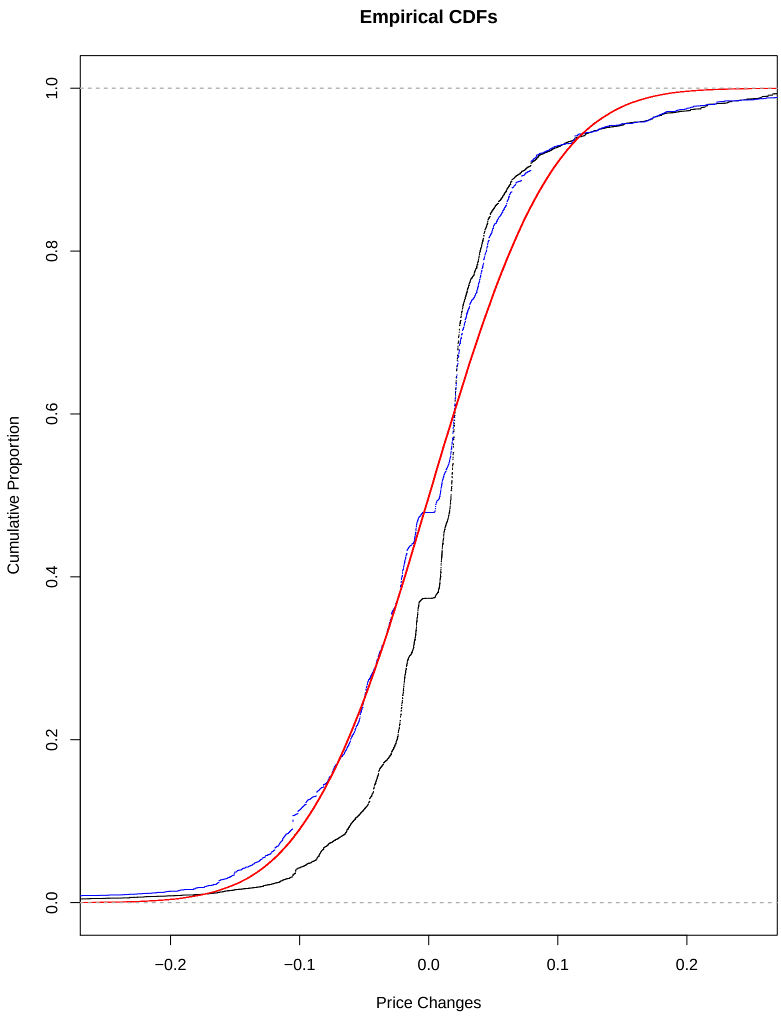

Figure 2: Size-Distribution of price updates (in log changes)

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

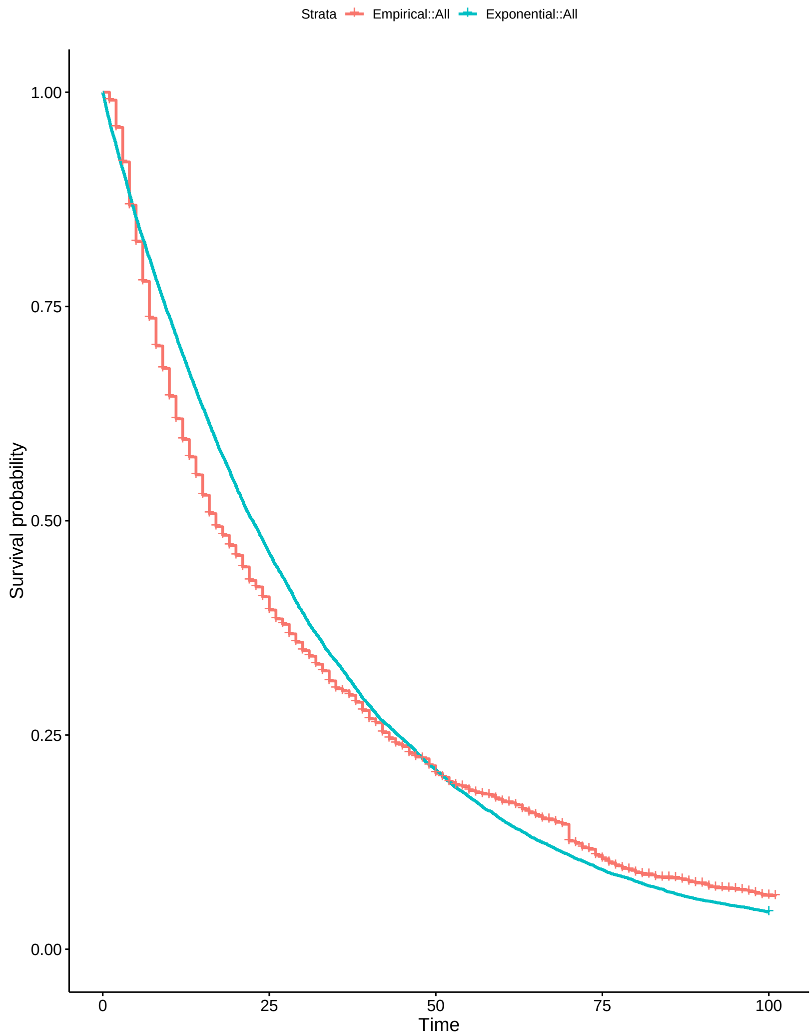

Figure 3: Distribution of arrival of updates

The red line plots the empirical survival function of price updates (Kepler-Meier estimate). The blue line depcits the theoretical exponential cdf with the same average arrival rate.

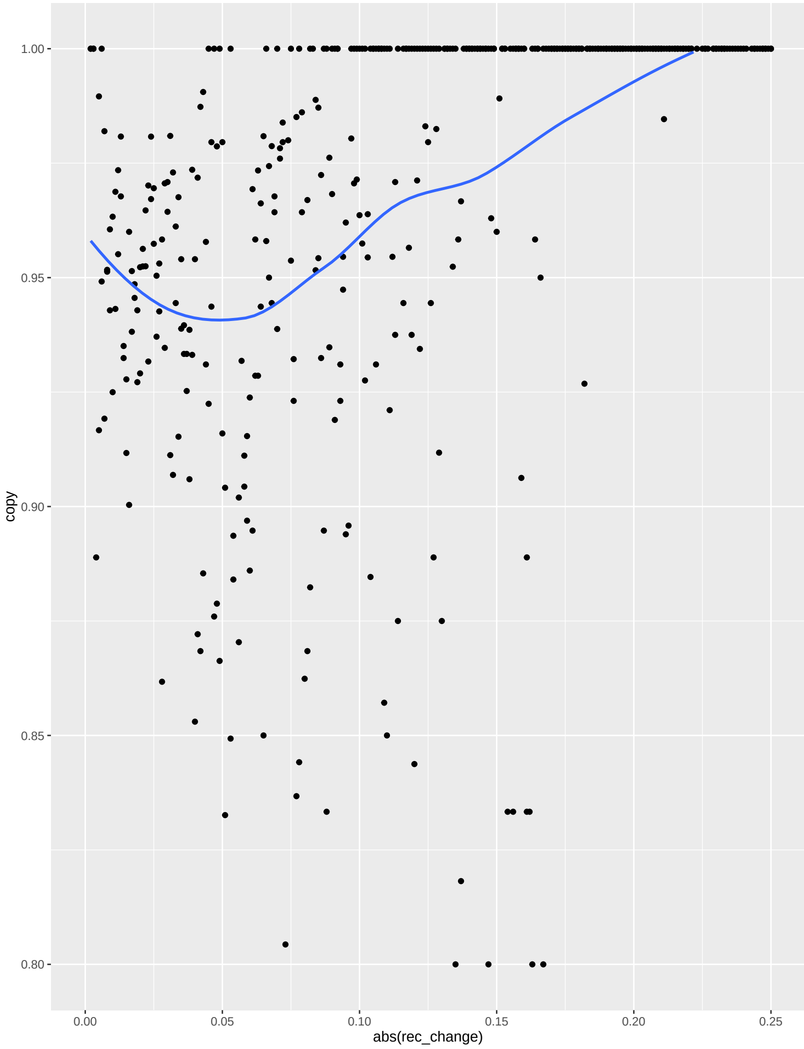

Figure 4: Matching the Recommendation

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.

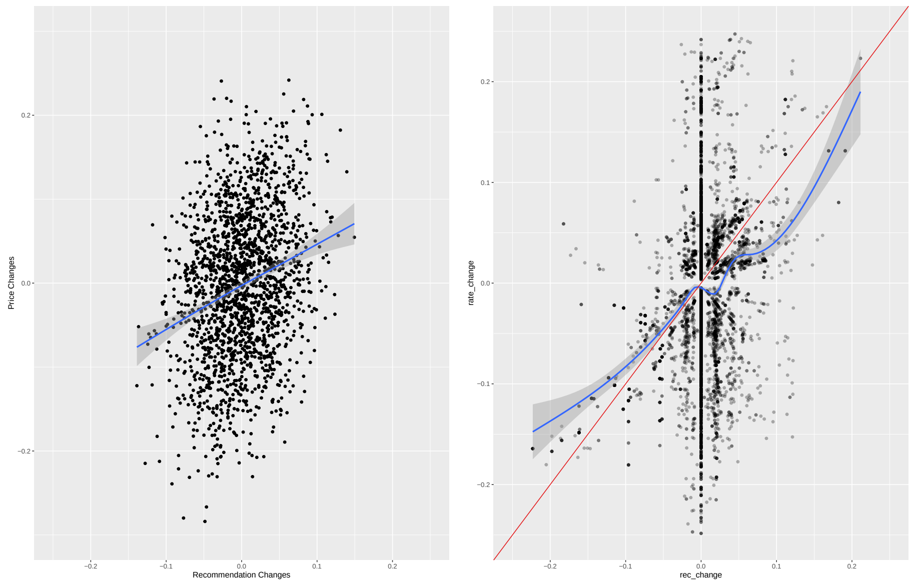

(a) Model

(b) Data

Figure 5: Model Fit: Recommendations and Prices

The left panel represents simulated data from the model. The right panel represents the true data. Each point represents an update that does not match the recommendation. The horizontal axis captures the log change in the recommendation. The vertical axis represents the change in the price.

Table 2: Pass-Through Rates of Recommendation

	Price Updates			
	All	Manually Updated		
Recommendation	0.974***	0.725***	0.733***	0.738***
	(0.002)	(0.005)	(0.006)	(0.006)
Days ahead Poly	No	No	Yes	Yes
Room × Month FE	No	No	No	Yes
N	76090	76090	76090	76090

Notes: The dependent variable is the cummulative change in the actual rate since the last update. Recommendation is the cummulative change in the recommendation since the last update. All regression include all updates; coefficients under *manually updated* are the interaction term (recommendations times manual). Significance levels: *** $p < 0.001$

Table 3: Targets for the Pooled Data

Moment	Data	Model
$\sqrt{V(p\|U)}$	0.068	0.069
$\sqrt{V(r\|U)}$	0.074	0.075
$\sqrt{V(p-r\|U)}$	0.035	0.035
$\sqrt{E(pr\|U)}$	0.068	0.067
$\Pr(C\|U)$	0.840	0.836
$\Pr(C\|U,L)$	0.947	0.958
$\Pr(U)$	0.038	0.038

Notes: The first two rows report the standard deviation of the price and the recommendation, both conditional on an update (U). The third row reports the standar 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), again conditional on an update. Rows five and six report the copy rate both unconditionally and conditional on the recommendation's change exceeding 10% (L). The last row reports the update rate.

Table 4: Estimates

	σ_x	σ_y	σ_z	σ_c	b_1	b_2
Pooled	0.034	0.045	0.056	1.86	0.035	0.543
A	0.030	0.103	0.088	1.29	0.0189	0.607
B	0.021	0.054	0.043	2.07	0.025	0.821
C	0.025	0.035	0.019	2.60	0.059	0.790
D	0.018	0.037	0.002	2.66	0.042	0.603
E	0.030	0.034	0.003	2.08	0.025	0.689
F	0.019	0.034	0.000	1.81	0.025	0.367
G	0.033	0.029	0.025	2.48	0.082	0.534
H	0.044	0.071	0.015	1.61	0.030	0.515
I	0.035	0.088	0.007	1.58	0.047	0.454

Notes: Estimated values for different hotels (A-I) and the pooled data.
Bootstrapped Standard errors to be reported.

Table 5: Counterfactuals

	Benchmark	Profit Loss	Delegation	Biased
A	0.995	0.990	0.954	0.960
B	0.992	0.983	0.914	0.921
C	0.984	0.970	0.717	0.750
D	0.984	0.968	0.809	0.831
E	0.980	0.961	0.671	0.710
F	0.996	0.989	0.760	0.826
G	0.980	0.962	0.574	0.631
H	0.983	0.961	0.731	0.775
I	0.990	0.974	0.862	0.864

Notes: The value in the first column corresponds to the welfare loss (profit plus costs) in the status quo relative to the profit loss under complete inaction. The value in the second column is the implied accounting profit loss (disregarding adjustment costs) relative to inaction. The third column represents the welfare loss in the counterfactual exercise (delegation) relative to inaction. The last column describes the welfare loss we would expect in a counterfactual where the algorithm continues to make biased recommendations (decisions) relative to inaction.

A Ommitted Proofs

Proof of Lemma 1. We first establish that $0 \leq l_1(r,c) \leq 2(1-\tilde{\lambda})^2 \int_{y \in Y_0(r,c)} d\Psi_y(y) r$, 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,

$$l_1(r,c) = 2(1-\tilde{\lambda}) \int_{y \in Y_0(r,c)} (r - \tilde{x}(r) - y) d\Psi_y(y). \tag{2}$$

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 (2) 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 $l_2(r,c) = b_2 \int_{y \notin Y_0(r,c)} d\Psi_y(y) > 0$. $\square$

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

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

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

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

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

$$\begin{aligned} 2\tilde{\lambda}^2 \bar{r} - l_1(\bar{r},c) &> 2\tilde{\lambda}^2 \bar{r} - 2\bar{r}(1-\tilde{\lambda})^2 \int_{y \in Y_0(\bar{r},c)} d\Psi_y(y) \\ &\geq 2\tilde{\lambda}^2 \bar{r} - 2(1-\tilde{\lambda})^2 \bar{r} \geq 0, \end{aligned}$$

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

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

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

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

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

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

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

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

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

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,

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

$$= \frac{1}{A}\int_0^{c(r)}\left(\int_{-\infty}^{r-\tilde{x}(r)-d(c)} y d\Psi_y(y) + \int_{r-\tilde{x}(r)+d(c)}^{\infty} y d\Psi_y(y)\right) dF(c), \qquad (3)$$

where

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

Notice then that,

$$\int_{-\infty}^{r-\tilde{x}(r)-d(c)} y d\Psi_y(y) + \int_{r-\tilde{x}(r)+d(c)}^{\infty} y d\Psi_y(y)$$

$$= \int_{-\infty}^{-(r-\tilde{x}(r))-d(c)} y d\Psi_y(y) + \int_{-(r-\tilde{x}(r))-d(c)}^{(r-\tilde{x}(r))-d(c)} ry d\Psi_y(y) + \int_{r-\tilde{x}(r)+d(c)}^{\infty} y d\Psi_y(y)$$

$$= \int_{-(r-\tilde{x}(r))-d(c)}^{r-\tilde{x}(r)-d(c)} y d\Psi_y(y) \leq 0 \qquad (4)$$

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 (3) must be negative. When $r < 0$ the inequality in (4) is simply reversed proving the proposition. □

Proof of Corollary 1. It is enough to show that

$$\frac{1}{A'}\int_{-\infty}^{\infty}\int_0^{c(r)}\int_{y\notin Y_0(c,r)} r(x) y d\Psi_y(y) dF(c) d\Psi_x(x) \leq 0, \qquad (5)$$

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 (4) when r is positive and reverses it when r is negative. Consequently, the inner double integral in (5) is always less than zero proving the corollary. □

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

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

Rewritting we have

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

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

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

$\square$