Strategic Responses to Algorithmic Recommendations: Evidence from Hotel Pricing

Daniel Garcia, Juha Tolvanen, Alexander Wagner

University of Vienna & University of Salzburg

EEA 24.08.2022

Motivation

- Organizations rely increasingly on **algorithmic systems** to make decisions.
    - Hiring and candidate selection (e.g. Kahn et al. 2019, Cowgill 2020)
    - Pricing (Assad et al. 2021, Garcia et al. forthcoming)
    - Bail decisions (Pro-Publica/Kleinberg et. al. 2019)
    - Allocation of public funds/efforts (e.g. Austrian Employment Services)
- In most of these instances, however, humans retain the final decision power: machines provide advice

Motivation

- Why do humans retain *formal* power?
  - Benefits:
    - Humans may have additional *soft* information that is not easily available to the algorithm
    - Humans may be able to *correct* Algorithmic Bias.
    - Humans are afraid/reluctant to AI decisions/errors (Dietvorst et al. 2015, Hidalgo et al. 2020).
  - Costs:
    - Human actions are plagued by errors, stereotypes and other biases.
    - Human decisions are **slow and costly** (time, effort)

What we do

- We model and estimate this trade-off in the context of hotel pricing in a number of hotels in Central Europe.
- Hoteliers receive price recommendations but retain final decision
  - Managerial inertia generates a conflict of interest and results in biased communication.
    - Lesson: Algorithm correcting decision maker's behavior can backfire
  - Simple model can generate patterns consistent with the data.
  - We structurally estimate the model to assess whether delegation to the algorithm would outperform recommendations.

Data

- Revenue management consultancy that provides price recommendations for 200+ hotels in Central Europe.
- We observe daily recommended rates and actual rates for a sample of 9 hotels over 10-12 months (5M+ observations)
- Prices are sticky: we see 130K+ price changes and 800K+ recommendation changes.

Data: More than Recommendations

- One important feature of our setting is that hoteliers find it easier to *accept* a recommendation than to choose their own price optimally.
- Hoteliers use a dashboard where they can see the current price and the recommended price and choose
  - Keep the current price (inaction).
  - Accept the recommendation.
  - Acquire further information (bookings, prices of competitors, etc.) and choose the price optimally.

4 Stylized facts

1. Prices stickier than recommendations
   - Prices change once every 35 days, recommendations every 7 days
2. The larger the recommended change, the higher the probability of a price update Update regs
   - Recommendations carry information about optimal prices
3. Probability of copying recommendation increases in the size of the recommended change Copy regs
   - Contrast to standard models of communication where advisor's influence decreases with *extreme* recommendations.
4. Conditional on *a manual update* only about 73% of the absolute recommended change gets passed into prices Pass-through regs
   - Hotelier seems to think that on average the recommender is exaggerating price increases and decreases.
   - Even if hotelier has private info, as long as interests are aligned, this should be 100%.

Model intuition: Moral Hazard

- Both, Algo and Hotelier, max hotel's profits - but only hotelier faces adjustment costs
- Hoteliers update infrequently since they face adjustment costs
- Too slow for Algo $\Rightarrow$ conflict of interest
- Algorithm has an incentive to *exaggerate* its signal to motivate adjustment.
- In equilibrium
  - If Hotelier puts effort, she can extract the signal and update perfectly.
  - If Hotelier relies on the recommendation, she ends up choosing a biased price.
  - Algo faces a trade-off: More exaggeration, more frequent adjustments but also worse prices when copying

Model in one pic

- Common loss function $L_h = (p - p^*)^2$. Rec's signal x, Rec$=r(x)$, hotelier's (costly) signals y, z, info- and adjustment costs c_1, c_2, and $p^* = x + y + z$.

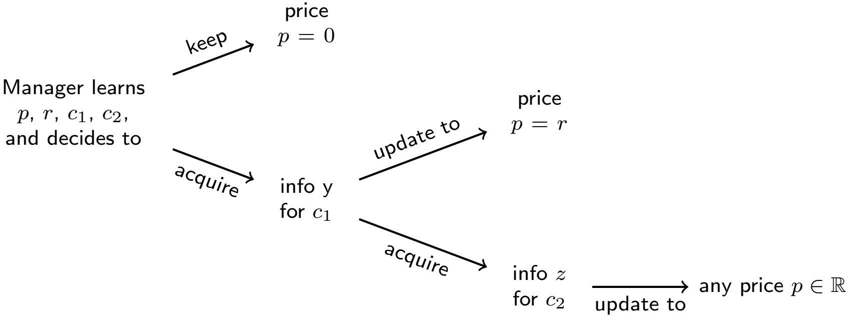

We look for linear eq: $r(x) = \alpha x$ for some $\alpha \in \mathbb{R}$, which hotelier correctly inverts in equilibrium. Parametrize and estimate with SMM.

Counterfatual: Delegation

Table: Counterfactuals

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

Details: All losses relative to complete inaction, Benchmark is current state More details

Conclusions

- We study algorithmic recommendations in a classical framework of advice.
- Private information held by human managers is sizeable but does not translate into superior decisions.
- Adjustment costs/inattention induce a natural conflict of interest, resulting in biased communication and decision-making.
- Delegation likely to improve outcomes, but gains vary widely across hotels.

Targets

Target	Data	Model
s.d. of r when update	0.068	0.069
s.d of p when update	0.074	0.075
(sq.root) cov. of p and r (update)	0.035	0.035
s.d. of $p - r$ when update	0.068	0.067
update rate	0.038	0.038
copy rate	0.840	0.836
copy rate (large r)	0.947	0.958

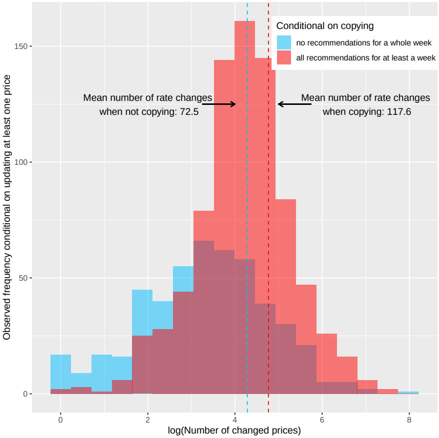

Figure: Log Price Changes By Group

Literature

- Economists have long studied the problem of **advice**
  - Cheap Talk models that emphasize nosiy communication (Crawford-Sobel 1982, Kartik et al. 2013)
  - Costly Information Transmission/Acquistion that emphasize the role of authority (Aghion-Tirole 1997)
- Applications to AI
  - Cowgill and Stevenson (2020) suggest a Crawford-Sobel model to understand human-machine interactions.
  - Agarwal et al. (2020) use an Aghion-Tirole model: human delegates *routine* decisions.
- Recent literature on humans as **evaluators** (Kleinberg et al (2018), Chan et al (2021), Mullainathan and Obermeyer (2022))

Probability of updating - regressions

Table: Price Update Probability

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

Notes: 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 AMnth refers to the arrival month. Significance levels: $^{***}p < 0.001$

Back

Probability of copying - regressions

Table: Price Update Copy Rates

	Copying Probability			
Rec Change	0.133^{***}	0.141^{***}	0.151^{***}	0.151^{***}
	(0.012)	(0.011)	(0.011)	(0.011)
Hotel × Date FE	No	Yes	No	No
Room × Date FE	No	No	Yes	No
R × D × AMnth 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 AMnth refers to the arrival month. Significance levels: *** $p < 0.001$

Back

Pass-Through of Recommendation - regressions

Table: Pass-Through Rates of Recommendation

	Change in actual price			
	All	Manually Updated		
Rec Change	0.974^{***}	0.725^{***}	0.733^{***}	0.738^{***}
	(0.002)	(0.005)	(0.006)	(0.006)
Days ahead Polynomial	No	No	Yes	Yes
Room $\times$ AMnth 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. Rec Change is the cumulative (log) change in the recommendation since the last price update. Coefficients for *Manually Updated* correspond to the interaction term of Rec Change $\times$ manual. *Room* is the room type and *AMnth* refers to the arrival month. Significance levels: *** $p < 0.001$

Back

Details on the counterfactuals

- Fhe first column (Benchmark) corresponds to the welfare loss in the status quo relative to the welfare loss under complete inaction
- The second 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

Back