ONLINE APPENDIX

EMILIO CALVANO
Università LUISS (Rome), Toulouse School of Economics and CEPR

GIACOMO CALZOLARI
European University Institute, Bologna University and CEPR

VINCENZO DENICOLÒ
Bologna University and CEPR

SERGIO PASTORELLO
Bologna University

The appendix offers additional results and detailed modeling information that were not included in the main text. To maintain coherence, we have preserved the section numbering to align with the corresponding sections in the main article. Consequently, Section A.x.y complements Section x.y in the main text and so forth.

3. MODEL-BASED RECOMMENDER SYSTEMS

3.1. *Estimation procedure: regularization and cross-validation*

In this section we discuss more details of the estimation procedure of the RS, its regularization and the cross-validation routine mentioned in Section 3.3 of the paper.

The minimization of the loss function (6) in the paper proceeds by calculating the gradient vector of partial derivatives w.r.t. the $J \times H + I \times H$ decision variables and using numerical methods to find improved values of $\hat{\mathbf{T}}, \hat{\mathbf{V}}$.

In vector notation, the gradient is

$$\mathbf{G} = -\begin{pmatrix} \mathbf{E}\hat{\mathbf{V}} - \lambda_t \hat{\mathbf{T}} \\ \mathbf{E}'\hat{\mathbf{T}} - \lambda_v \hat{\mathbf{V}} \end{pmatrix} \tag{1}$$

where $\mathbf{E}$ is the $I \times J$ matrix of residuals $e_{ij} = \tilde{r}_{i,j} - \sum_{h=1}^{H} \hat{t}_{ih}\hat{v}_{jh}$.

The estimation algorithm is a version of the Stochastic Gradient Descent algorithm that is available upon request. Given some arbitrary initial values for $\hat{\mathbf{T}}_0, \hat{\mathbf{V}}_0$, at iteration t the algorithm computes an estimate of $\mathbf{G}_t$ and iterates according to:

$$\begin{pmatrix} \mathbf{T}_{t+1} \\ \mathbf{V}_{t+1} \end{pmatrix} = \begin{pmatrix} \mathbf{T}_t \\ \mathbf{V}_t \end{pmatrix} - \eta \mathbf{G}_t, \tag{2}$$

where $\eta > 0$ is the step size and $\mathbf{T}_{t+1}, \mathbf{V}_{t+1}$ are the new values of the variables of interest at the end of the iteration. The procedure is stopped at 50 iterations, to avoid overfitting and reduce the computational burden.

Emilio Calvano: ecalvano@luiss.it
Giacomo Calzolari: giacomo.calzolari@eui.eu
Vincenzo Denicolò: vincenzo.denicolo@unibo.it
Sergio Pastorello: sergio.pastorello@unibo.it

The regularization parameters λ_t, λ_v and η and the number of factors H when not pre-specified (more on this below), are determined with a 5-fold cross-validation, which is customary in machine learning.

In particular, we partition the ratings into five equally sized subsets. We then use four of these sets to estimate the model and compute the out-of-sample prediction error for the fifth set, obtaining a measure of the predictive ability of the model (RMSE). For a given parameterization, this procedure is replicated changing the left-out set, and the average RMSE is computed.

The entire procedure is then replicated for a range of possible values of λ_t, λ_v and η (as well as H when it is not pre-specified). Specifically, we have used the grid $\lambda_t, \lambda_v, \eta \in \{0.01, 0.05, 0.1, 0.15, 0.20\}$. We then choose the values of the hyper-parameters that minimize the average RMSE as computed above.

In one of the robustness checks, we allow the algorithm to choose the number of latent factors by means of a cross-validation procedure (results are reported in supplementary material). In this case, holding the true number of latent factors equal to $H = 2$, we let the algorithm choose the estimate $\hat{H}$ from a grid $1, 2, 3, 4, 5$[1]. Table 1, which reports the fraction of sessions where the number of latent factors is estimated at each of the possible five values, shows that the cross-validation routine almost systematically identifies the correct number of factors, $H = 2$. In less than 5% of the cases, the algorithm chooses $\hat{H} = 3$ instead.

H/α	0	0.25	0.5	0.75	1
1	0.00	0.00	0.00	0.00	0.00
2	0.95	0.96	0.95	0.96	0.97
3	0.05	0.04	0.05	0.04	0.03
4	0.00	0.00	0.00	0.00	0.00
5	0.00	0.00	0.00	0.00	0.00

TABLE 1. Each cell reports the fraction of cases in which the cross-validation procedure identifies a number of factors H as indicated in each rows, for different values of α (in the columns).

4. PRODUCTS AND PREFERENCES

4.1. *Preferences and transportation costs*

As mentioned in the main text, our model of horizontal product differentiation shares certain similarities with the conventional Hotelling framework. To emphasize these resemblances, we calculate the "transportation-cost" implicit in our framework when $\alpha = 0$ and compare it to the cost associated with more traditional Hotelling frameworks.

To calculate the implicit transportation cost, we employ polar coordinates as in Perego (2020), denoting product positions as $\{t_{i1} = \cos\theta, t_{i2} = \sin\theta\}$ and consumer positions as $\{v_{j1} = \cos\upsilon, v_{j2} = \sin\upsilon\}$. Using these coordinates, a product's location is represented by υ within the range $\left[0, \frac{\pi}{2}\right]$, and a consumer's location is represented by θ within the same range. Therefore, $\theta - \upsilon$ represents the angular distance between a consumer located at θ and a product located at υ. The utility that a consumer at θ derives from a product at υ is given by:

$$u(\theta, \upsilon) = \cos\theta\cos\upsilon + \sin\theta\sin\upsilon = \cos(\theta - \upsilon). \tag{3}$$

[1] In another robustness check we set the number of latent factors to $H = 5$. In this case we assume the algorithm knows the true number of latent factors. Results are in section **??**.

Since the utility that consumers receive from their ideal product (where $\upsilon = \theta$) is 1, the "transportation cost" is simply $1 - \cos(\theta - \upsilon)$; there is no explicit transportation cost parameter.

Figure 1 contrasts our transportation cost with a quadratic cost.

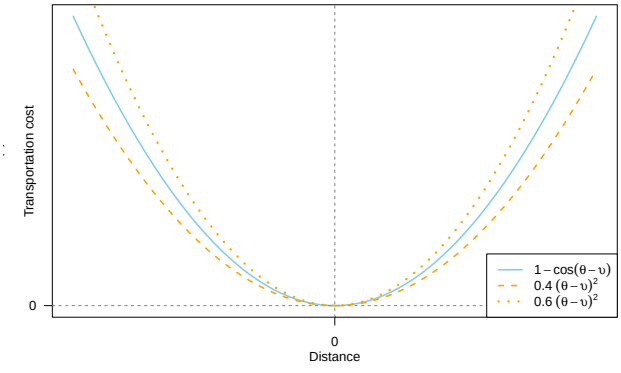

FIGURE 1.—The continuous curve depicts the transportation cost $1 - \cos(\theta - \upsilon)$ in our model. For comparison, the figure also presents the standard quadratic transportation cost $\tau(\theta - \upsilon)^2$. The dashed curve corresponds to $\tau = 0.4$, and the dotted one to $\tau = 0.6$. For $\tau = 0.5$, the quadratic transportation costs are almost indistinguishable from that of our latent factor model.

6. MARKET CONCENTRATION

6.3. *The superstar effect with horizontal differentiation*

Table IV shows that the superstar effect is significant when vertical differentiation is important (large α), but not when product differentiation is mainly horizontal (small α). In fact, RS creates superstars even under horizontal product differentiation. This tendency is not apparent in Table IV because with $m_{\mathcal{M}} = 19$ products there are many products that are "central," and the ones selected by the RS as superstars vary randomly from session to session. To eliminate this effect, we rank products in terms of their market share and calculate the average share, across the 100 sessions, of the most popular product, the second most popular, and so on. We do so for the case with RS and for the benchmark of unassisted search. The resulting distribution is depicted in Figure 2.

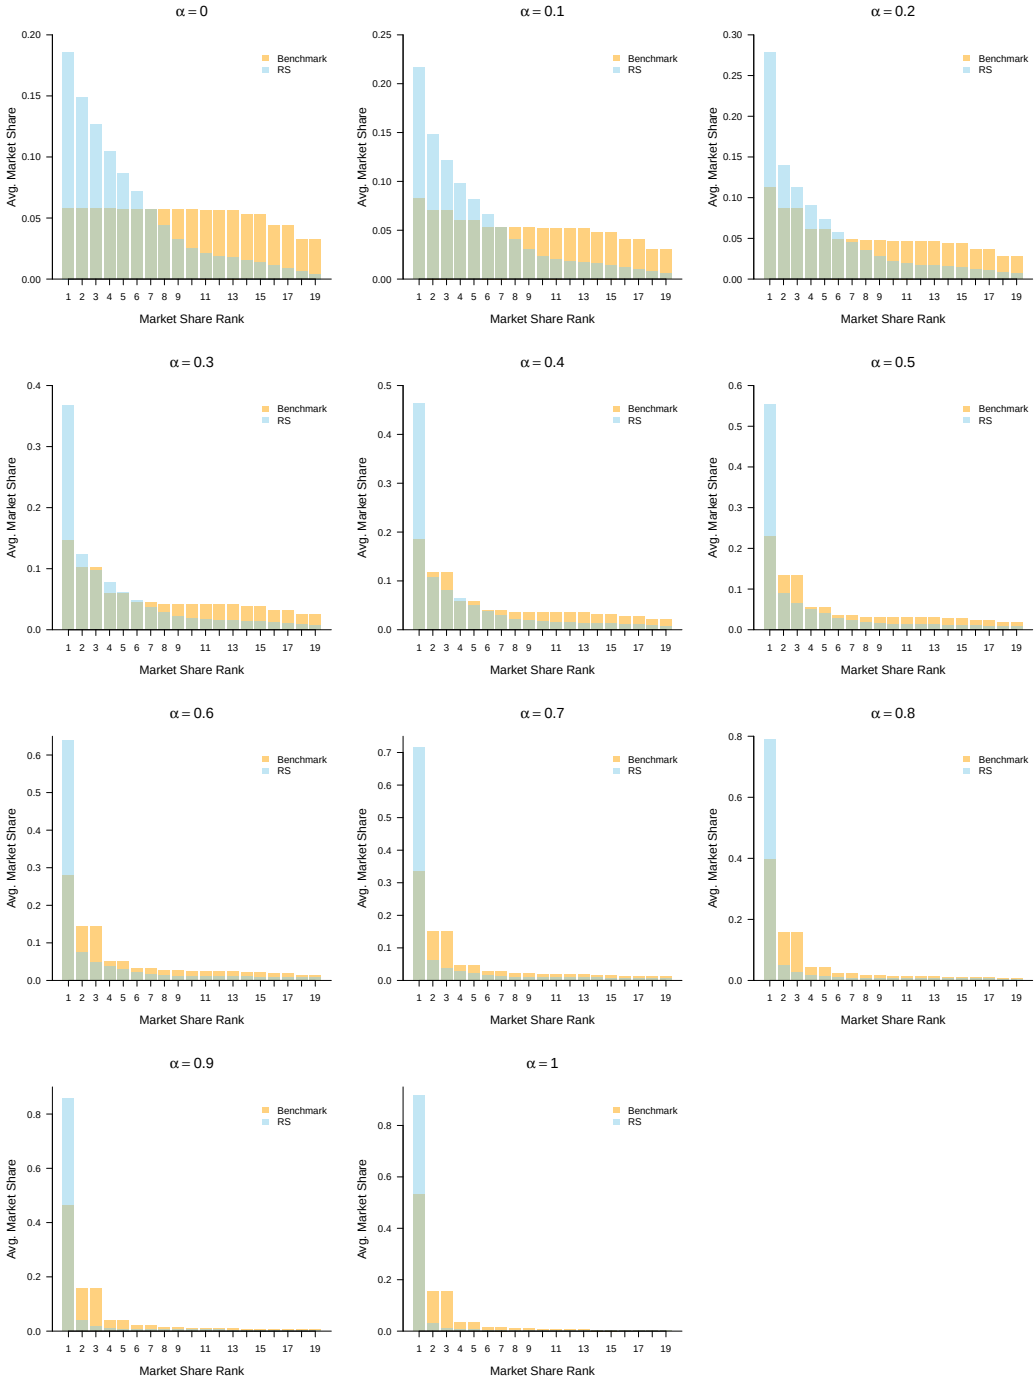

FIGURE 2.—Average Market share of products ordered by popularity, for different degrees of product differentiation ranging from pure horizontal ($\alpha = 0$), and pure vertical ($\alpha = 1$). The average is computed across session.

In the individual search benchmark, all products except for the most peripheral ones have similar market shares of around 5% when $\alpha = 0$. However, with the RS, the most popular product has an average market share of over 15%, indicating the presence of a significant superstar effect even when $\alpha = 0$. This effect becomes even stronger as α increases.

Reconsidering the definition of a 'superstar' as the product with the highest market share, we can identify this product in Figure 3. The figure illustrates the frequency with which each product attains superstar status across 100 simulations. Products are organized based on their position within the product space, spanning from the x- to the y-axis, as previously shown in Figure 1. Notably, central products exhibit a greater tendency to become superstars even in cases where $\alpha = 0$. As α increases, a noteworthy trend emerges: the product located at the center of the product space (characterized by $m_{\mathcal{M}} = 19$ and product number 10) is selected more frequently. In fact, for values of α equal to or exceeding 0.3, this specific product consistently assumes the role of the superstar.

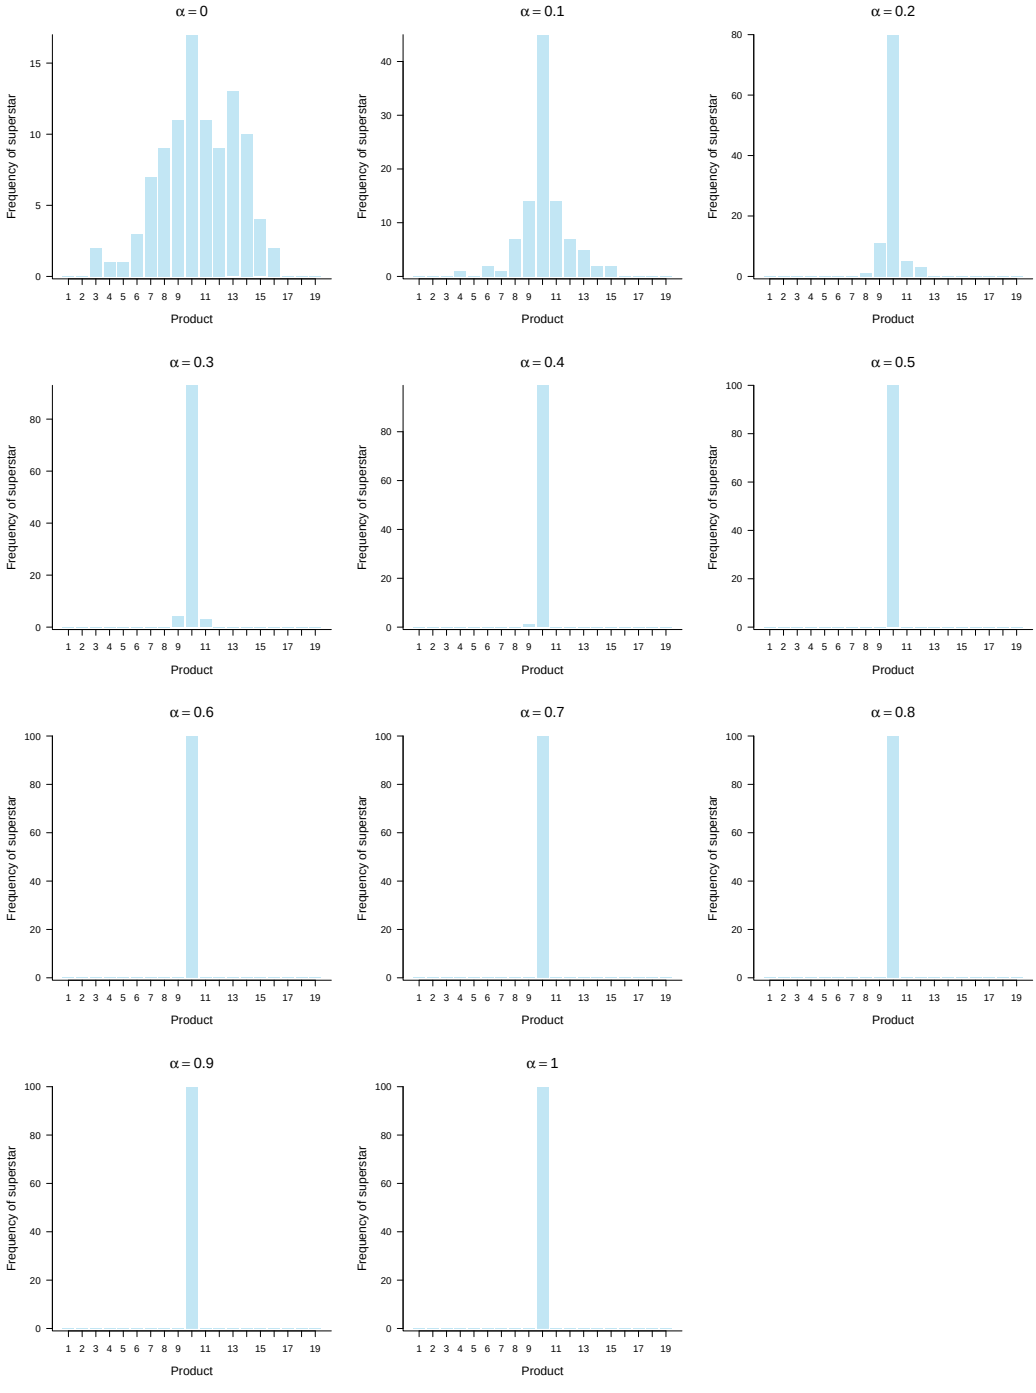

FIGURE 3.—Frequency with which each product is selected as the superstar. Products are ordered on the basis of their location in Figure 1 (body of the paper), from the x-axis (product 1) to the y-axis (product 19). The central product is product 10.

6.4. *The uniformity effect*

In Figure 4, we study the uniformity effect for different values of the density d of the rating matrix. In particular, in examining the estimation biases across varying values of d, a noteworthy pattern emerges. Except for the lowest value of d, the biases appear to be consistently similar for θ, β, and the sum of squares of β. This consistent behavior suggests that the estimation process remains stable for these parameters over a range of d values. However, at the lowest d, instead of observing increasing heterogeneity in the estimated β, as seen in all other cases, there is a homogenization of estimates of β. This homogenization aligns with what is consistently observed for estimates of θ. Furthermore, for this lowest d, there's a bias in favor of central products. In contrast, for all other d values, there's a bias leaning towards niche products.

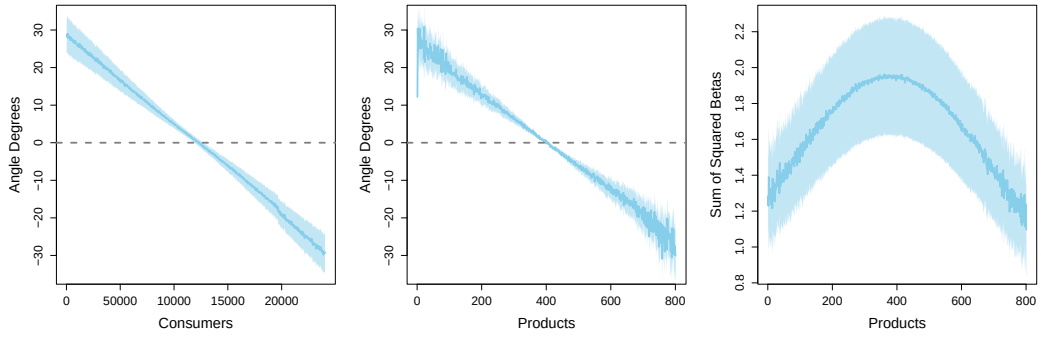

(a) Estimation biases with density $d = 0.6\%$.

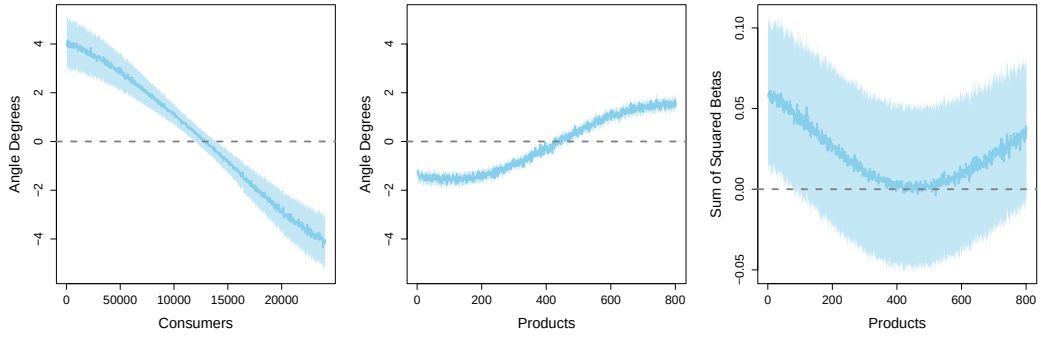

(b) Estimation biases with density $d = 1.2\%$.

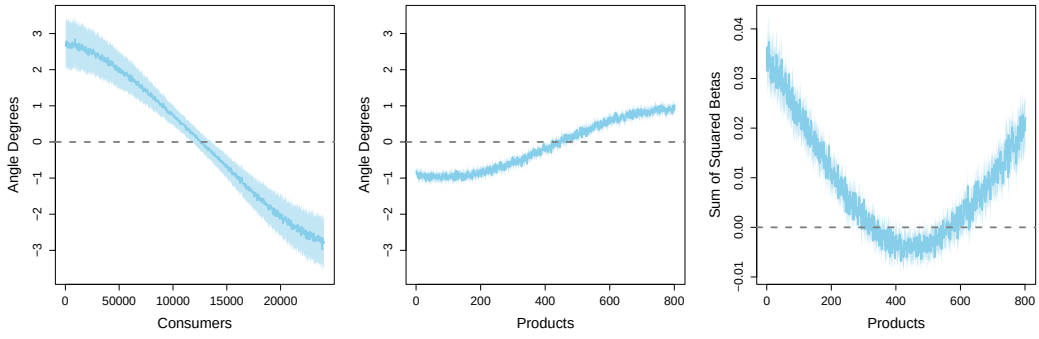

(c) Estimation biases with density $d = 2.4\%$.

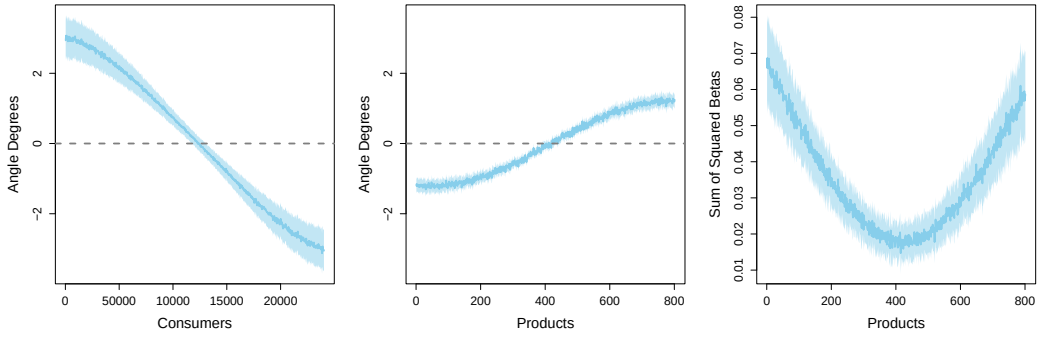

(d) Estimation biases with density $d = 3.6\%$.

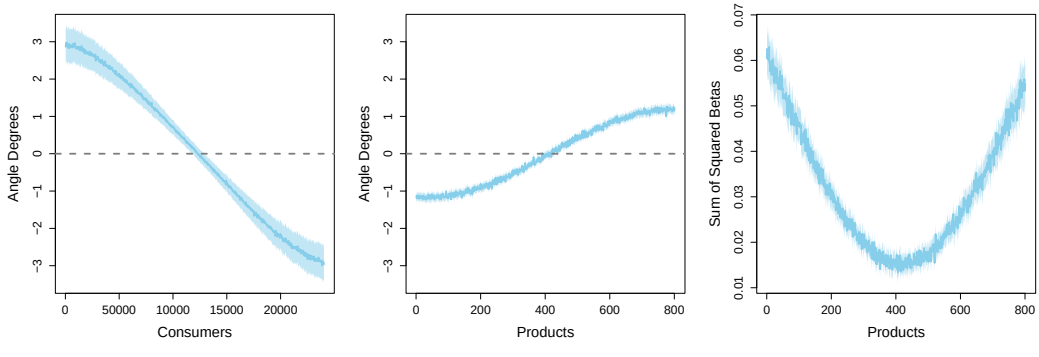

(e) Estimation biases with density $d = 4.8\%$.

FIGURE 4.—Estimation biases for various densities.

7. EQUILIBRIUM PRICES

7.1. *Deriving demand functions*

The demand functions are calculated as follows. Let q_{ij} denote the probability that the surplus provided by product j to consumer i surpasses the cutoff $\tilde{u}_{ij} - p_j$. The probability x_{ij} that consumer i purchases product j then is:

$$x_{ij} = \frac{1}{n_{\mathcal{M}}} q_{ij} + \left[\sum_{l \in \mathcal{M}} \frac{1}{n_{\mathcal{M}}} (1 - q_{il}) \right] x_{ij} \tag{4}$$

The first term on the right hand side is the probability that product j is sampled first and is immediately purchased. The term inside square brackets is the probability that no product is purchased after the first search, in which case the search process starts anew. Solving this equation yields:

$$x_{ij} = \frac{q_{ij}}{\sum_{\ell \in \mathcal{M}} q_{i\ell}}. \tag{5}$$

The total demand for product j is obtained by aggregating individual demands x_{ij} across consumers. Since q_{ij} decreases with p_j, this yields a system of decreasing demand curves.

With RSs, the algorithm recommends to consumer i the product

$$j^\star(i, \mathcal{M}) = \arg\max_{j \in \mathcal{M}} \hat{r}_{ij}. \tag{6}$$

The demand for this product is:

$$x_{ij^\star(i,\mathcal{M})} = q_{ij^\star} + (1 - q_{ij^\star}) \frac{q_{ij^\star}}{\sum_{\ell \in \mathcal{M}} q_{i\ell}}. \tag{7}$$

whereas the demand for products $j \neq j^\star$ is

$$x_{ij} = (1 - q_{ij^\star}) \frac{q_{ij}}{\sum_{\ell \in \mathcal{M}} q_{i\ell}}. \tag{8}$$

7.2. *Individual profits*

In section 7 of the main text, we mention the effect of the RS on equilibrium profits. Figures 5 and 6 illustrate these results in detail, showing how the RS affects the profits earned by each of the 19 firms in the market. The figure shows that when $\alpha = 0$, the peripheral firms definitely lose, while the gains and losses for the other firms do not seem to follow a clear pattern. When $\alpha > 0$, on the other hand, the central firm emerges as the clear winner, while the losers include not only the peripheral firms but also the closest competitors of the central firm. For the remaining firms, the effect is not significantly different from zero. In the figure, the median profit of central products decreases because creating champions with highly asymmetric market shares, the RS sometimes benefits and other times disadvantages those products. However, as anticipated from the price hikes, the average profits significantly increase, as show in Table 2.

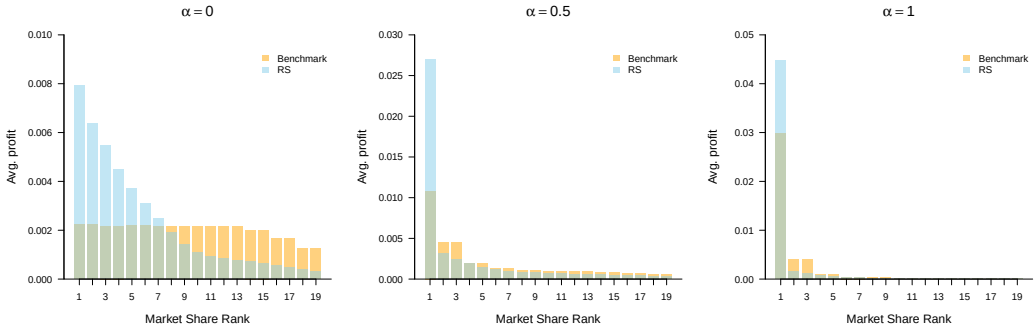

FIGURE 5.—The average (across sessions) profits of the suppliers ranked by market share.

α	0	0.25	0.5	0.75	1
*Un*assisted Search	0.0022 (0.0000)	0.0053 (0.0000)	0.0108 (0.0000)	0.0192 (0.0000)	0.0298 (0.0000)
*R*ecommender *S*ystem	0.0031 (0.0003)	0.0148 (0.0004)	0.027 (0.0002)	0.0372 (0.0001)	0.0448 (0.0001)
$\frac{RS-Un}{Un} \times 100$	40.91%	179.25%	150.00%	93.75%	50.34%

TABLE 2

AVERAGE PROFITS

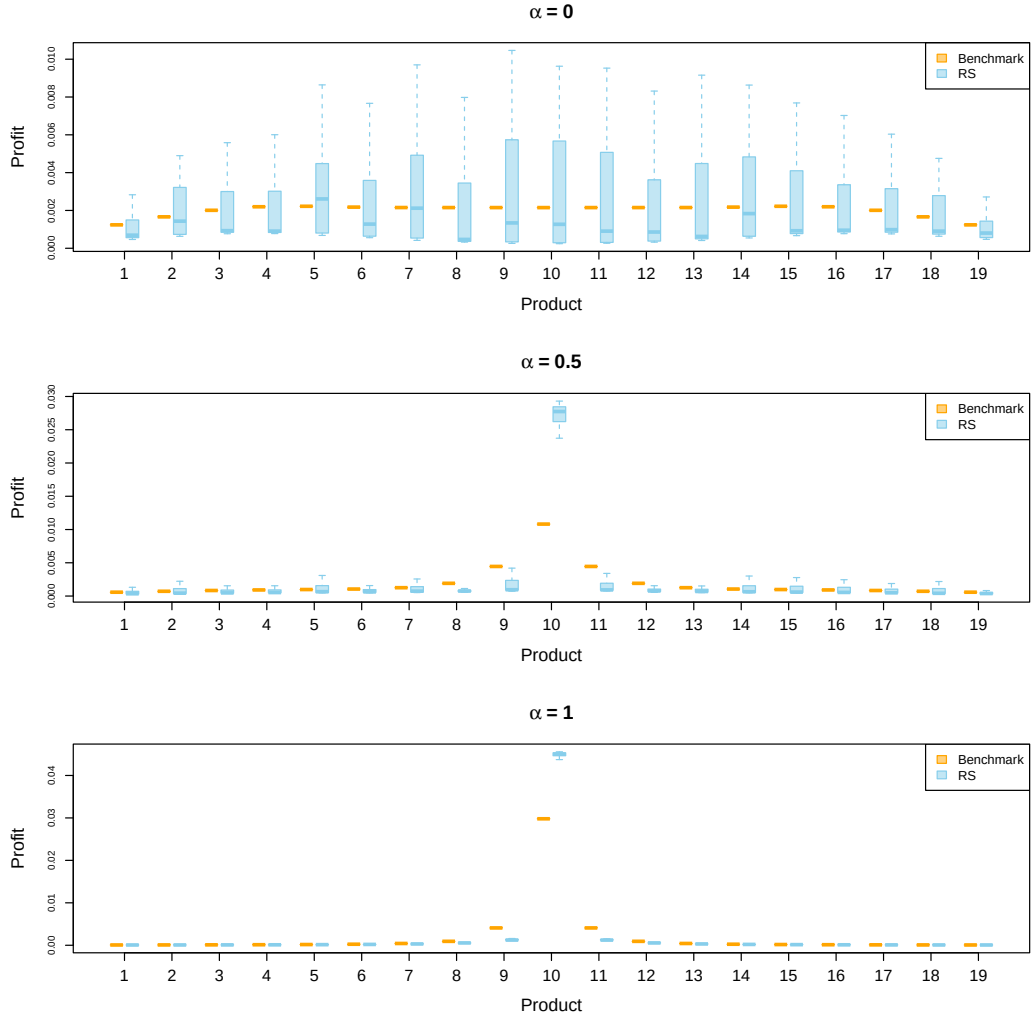

FIGURE 6.—The profits of the suppliers of the different products, ranked from 1 to 19 according to their location in the space of products: products 1 and 19 are peripheral, product 10 is central. As customary, the figure shows median, inter-quartile range (IQR), and the values obtained by subtracting 1.5 times the IQR range from the first quartile (Q1), and obtained by adding 1.5 times the IQR to the third quartile (Q3). That is the highest and lowest value excluding outliers).

7.3. *Number of products and search cost*

In this section, we illustrate how our results change when we vary the number of products $m_{\mathcal{M}}$ and the level of the search cost c_s. We consider both the effect on market concentration when prices are exogeneously set at zero, and the effect on equilibrium prices and consumer surplus.

Specifically, we consider three possible levels of the search cost: $c_s = 0.002$, $c_s = 0.004$ (the baseline value considered in the main text), and $c_s = 0.006$. For each of these values, and for each variable of interest, we present a heat map obtained by varying the number of products $m_{\mathcal{M}}$ from 7 to 31 (the baseline value considered in the main text being $m_{\mathcal{M}} = 19$), and the degree of vertical product differentiation α from 0 to 1. The heat maps represent the percentage change in the relevant variables with RSs relative to the individual search benchmark.

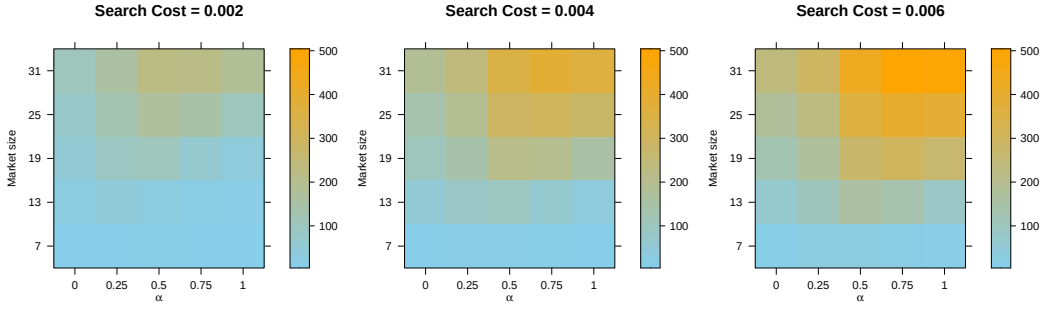

FIGURE 7.—Percentage shift of the HHI Index from Benchmark to RS.

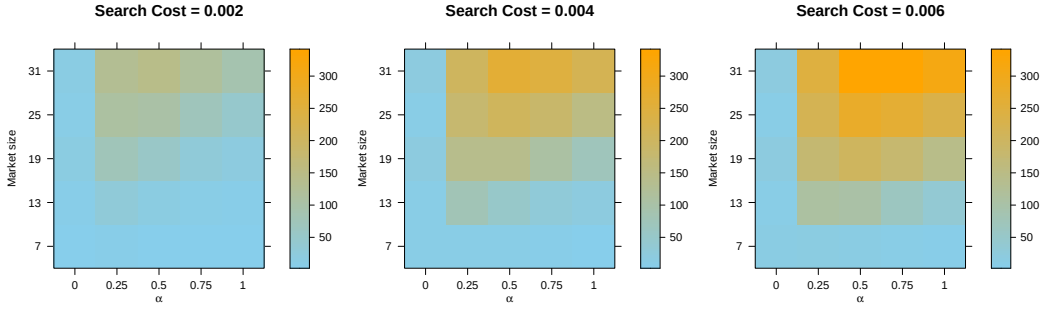

FIGURE 8.—Percentage shift of the market share of the central product from Benchmark to RS.

We have seen in the main text that the impact of the RS on market concentration, as measured by the HHI, increases with α. Figure 7 confirms this finding and shows that the effect also increases with the number of products and the level of the search cost. Unsurprisingly, the superstar effect (Figure 8) and the reverse long-tail effect also exhibit a similar pattern.

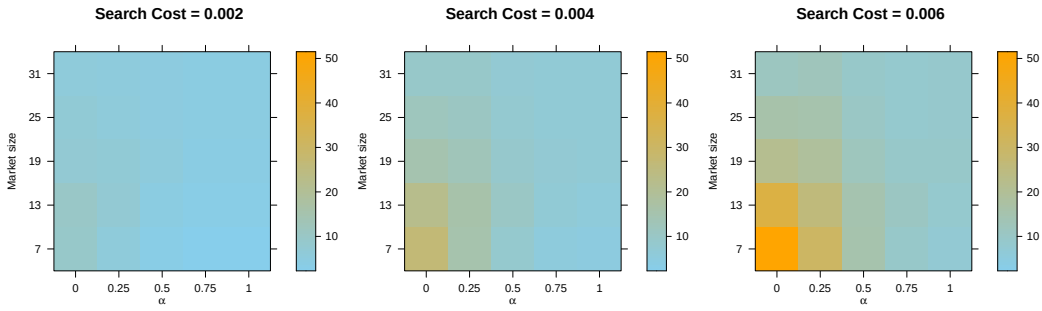

FIGURE 9.—Percentage shift of equilibrium prices (average across products) from benchmark to RS

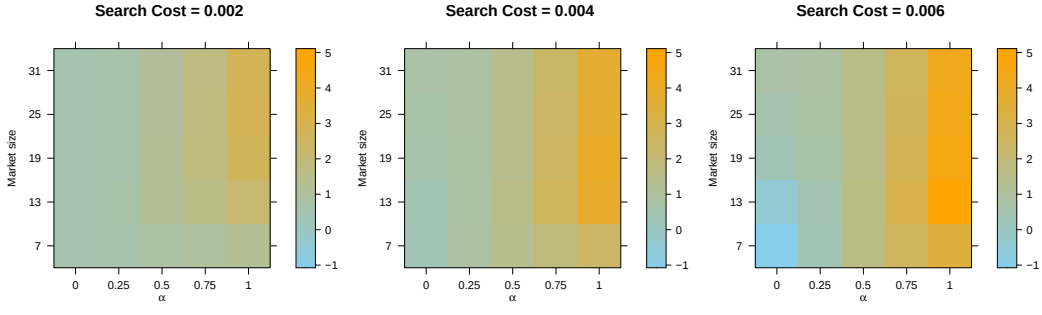

FIGURE 10.—Percentage shift of consumer surplus net of search costs from benchmark to RS.

In the main text, we have seen that the effect of the RS on equilibrium prices is highest when product differentiation is mainly horizontal (small α). Again, Figure x confirms this finding. Furthermore, it shows that the increase in equilibrium prices increases with the search cost and decreases with the number of products. When the search cost is high and the number of products is small, the average price may increase by as much as 50%.

In the paper, we observed that the RS has the most significant impact on equilibrium prices when product differentiation is primarily horizontal. Figure 9 reaffirms this. Additionally, the figure illustrates that the rise in equilibrium prices is proportional to the search cost and inversely related to the number of products. In cases where the search cost is high and the number of products is limited, the average price may experience an increase of up to 50

Figure 10 shows that this large increase in the average price may more than offset the positive effect of the RS on the matching of consumers and products, lading to a decrease in the level of consumer surplus. This possibility arises when the search cost is high ($c_s = 0.006$), product differentiation is horizontal ($\alpha = 0$), and the number of products is small ($m_{\mathcal{M}} = 7$ or 13).

9. BIASED RECOMMENDATIONS

In this section we provide further information for the case in which the RS provides manipulated recommendations, as discussed in the main text.

9.1. *Intensity of competition*

We will now dissect the average impact of manipulation on prices, illustrating its effects on various products while varying the favored product's identity. Figures 11-13 complement the main paper's findings by presenting results with a 50% manipulation rate, where the central product is favored, for different values of α.

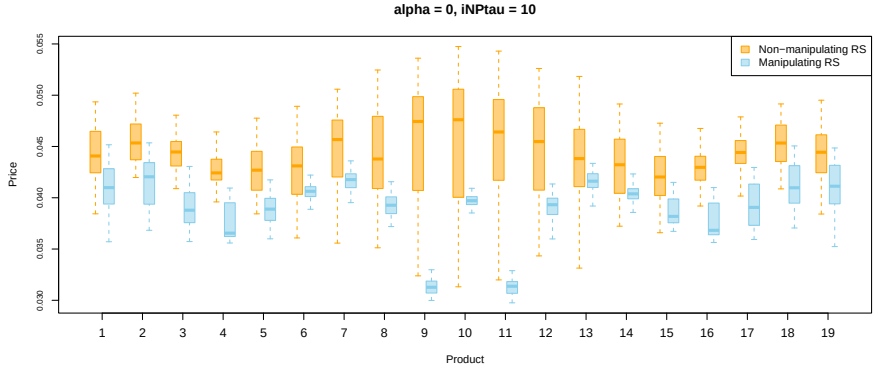

FIGURE 11.—Impact of 50% manipulation on prices for all products when product 10 is favored, with $\alpha = 0$.

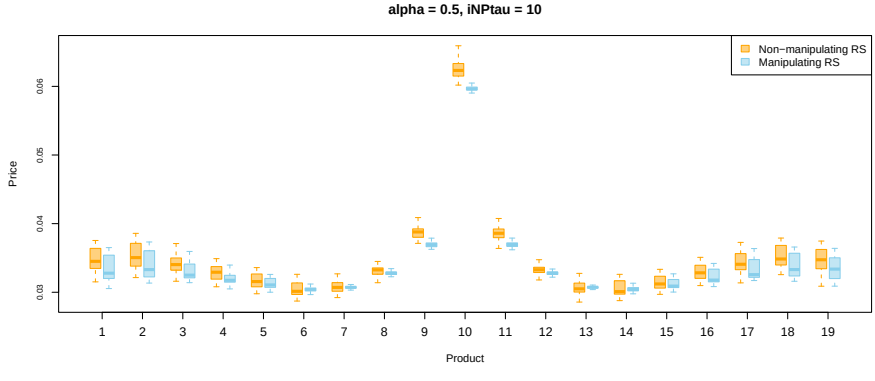

FIGURE 12.—Impact of 50% manipulation on prices for all products when product 10 is favored, with $\alpha = 0.5$.

The next figures 14, 15, 16 consider the same analysis of the effects of manipulation on prices, this time varying the favored product's identity. In all cases, manipulation intensifies competition and reduces equilibrium prices.

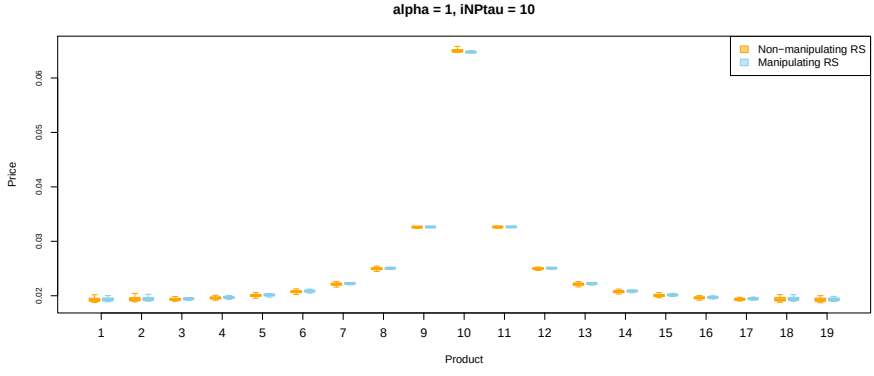

FIGURE 13.—Impact of 50% manipulation on prices for all products when product 10 is favored, with $\alpha = 1$.

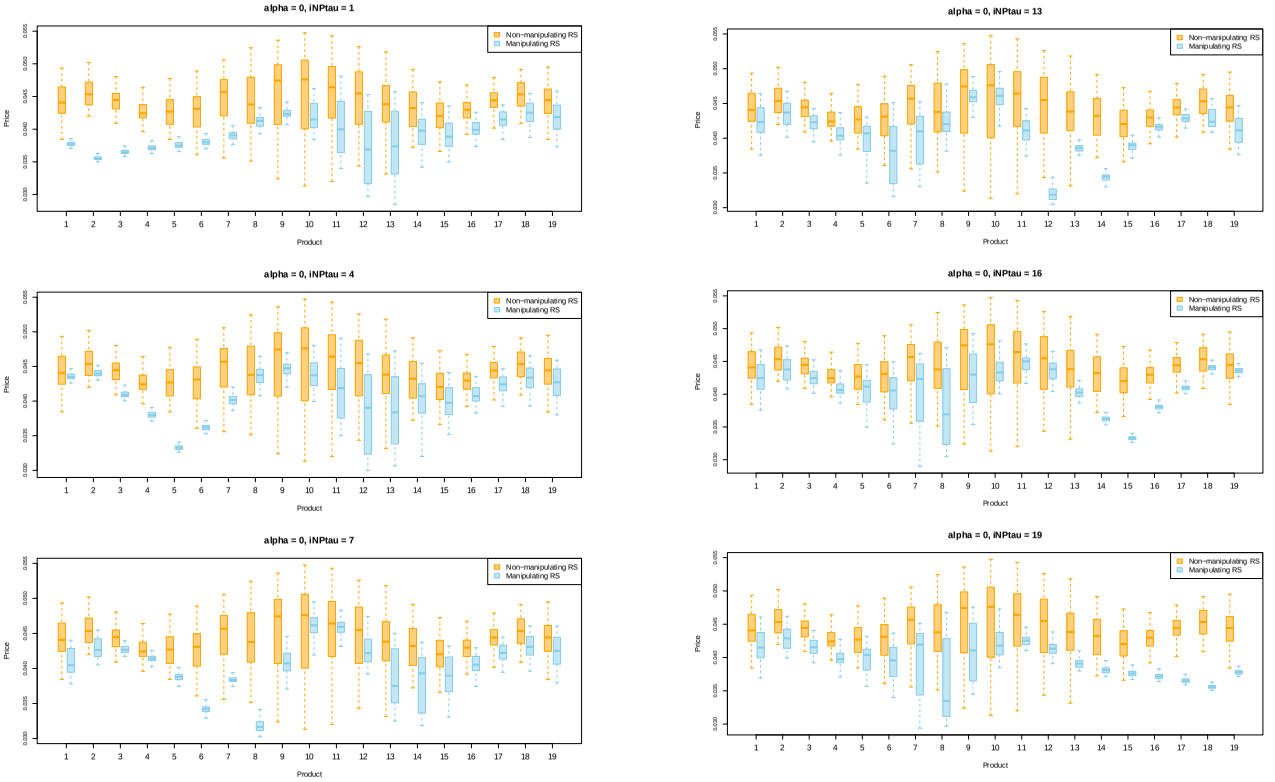

(a) (b)

FIGURE 14.—Effect of manipulation on equilibrium prices, when the favored product differs from the central one.'INPtau' is the favored product's identity. Our labeling convention designates products based on their location, with products 1 and 19 as peripheral and product 10 as central. This figure: parameters $\alpha = 0$, and manipulation frequency maintained at 50%.

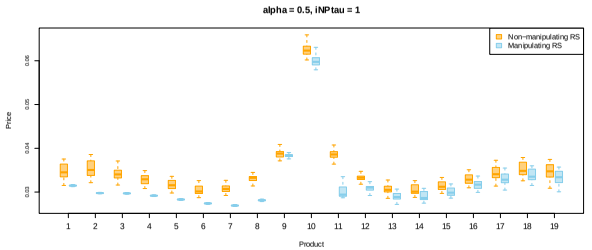

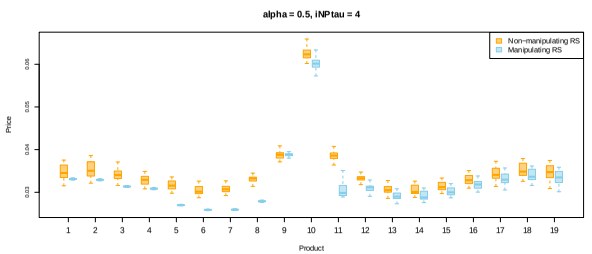

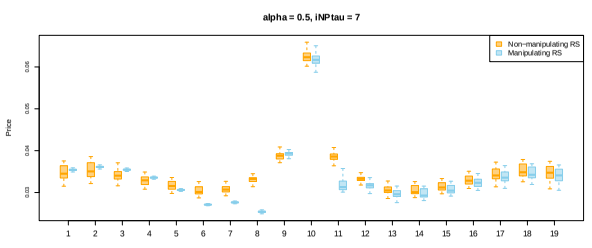

(a)

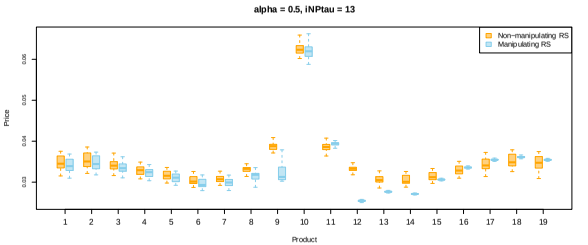

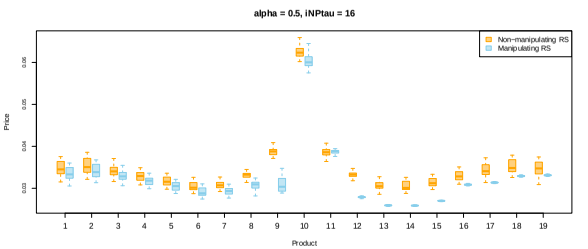

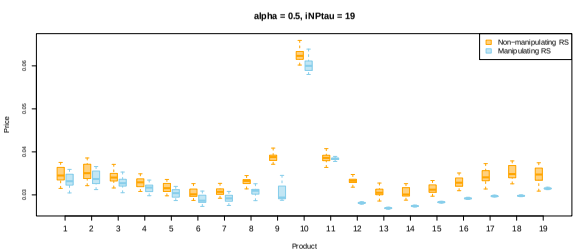

(b)

FIGURE 15.—Effect of manipulation on equilibrium prices, when the favored product differs from the central one.'INPtau' is the favored product's identity. Our labeling convention designates products based on their location, with products 1 and 19 as peripheral and product 10 as central. Key this figure: parameters $\alpha = 0.5$, and manipulation frequency maintained at 50%.

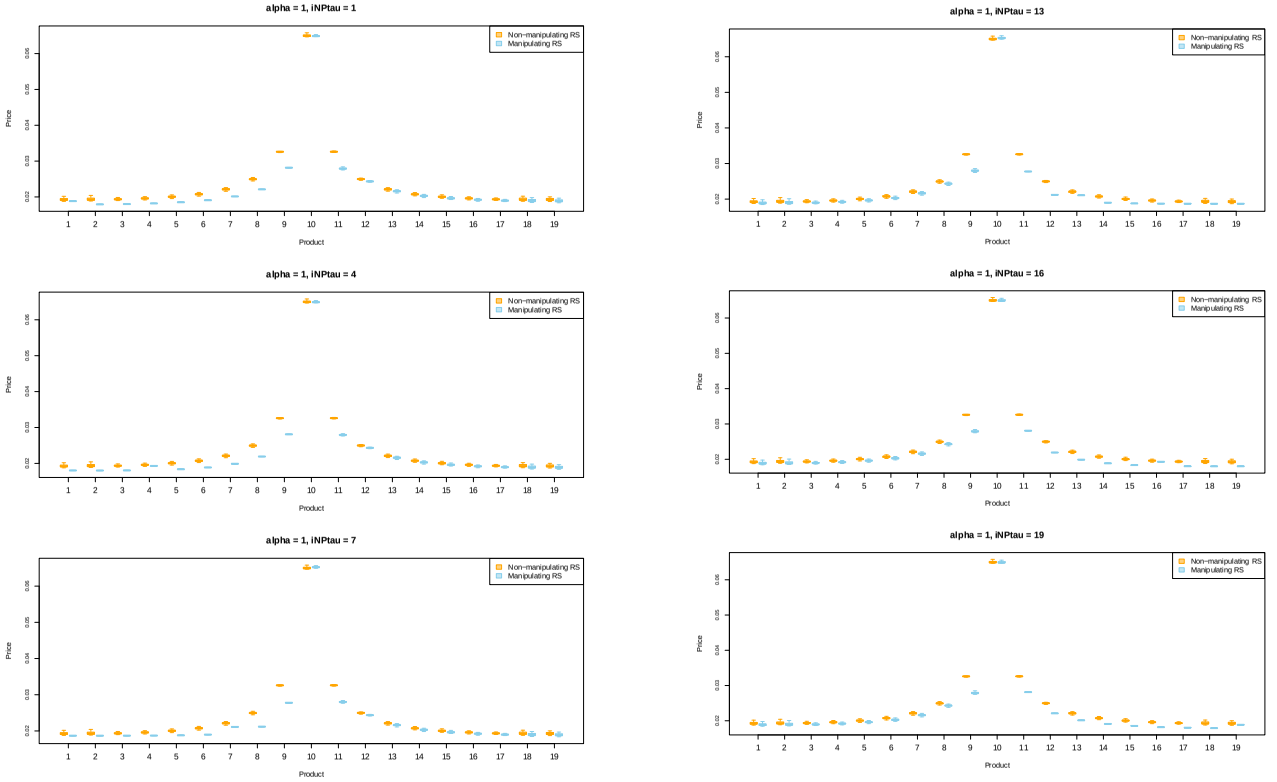

(a) (b)

FIGURE 16.—Effect of manipulation on equilibrium prices, when the favored product differs from the central one.'INPtau' is the favored product's identity. Our labeling convention designates products based on their location, with products 1 and 19 as peripheral and product 10 as central. Key this figure: parameters $\alpha = 1$, and manipulation frequency maintained at 50%.

9.2. *Profits*

We here unpack the effect of manipulation on profits by showing the impact on the different products and by varying the identity of the favored product. Figure 17 complements the main paper's findings by presenting results with a 50% manipulation rate, where the central product is favored. Manipulation has a strong positive effect on the favorite product and a negative effect on neighboring products that diminishes with the distance.

In analogy with the analysis on prices, we repeat the exercise by changing the identity of the favored product (figures 20, 21 and 22) In all instances, manipulation heightens competition and diminishes equilibrium profits.

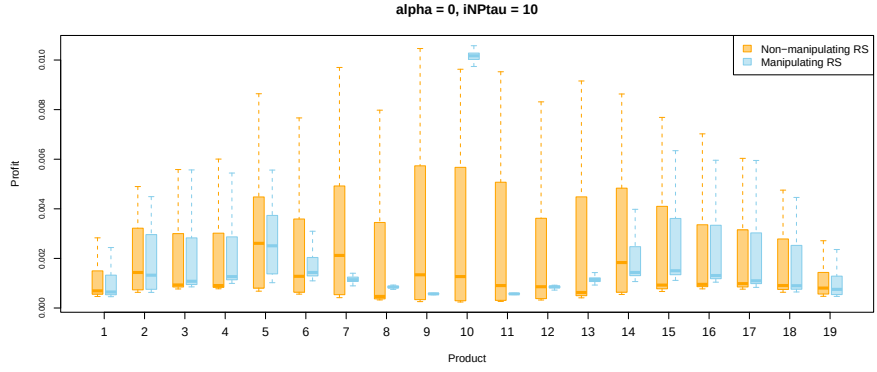

FIGURE 17.—Impact of 50% manipulation on profits for all products with product 10 favored by the recommender system ($\alpha = 0$).

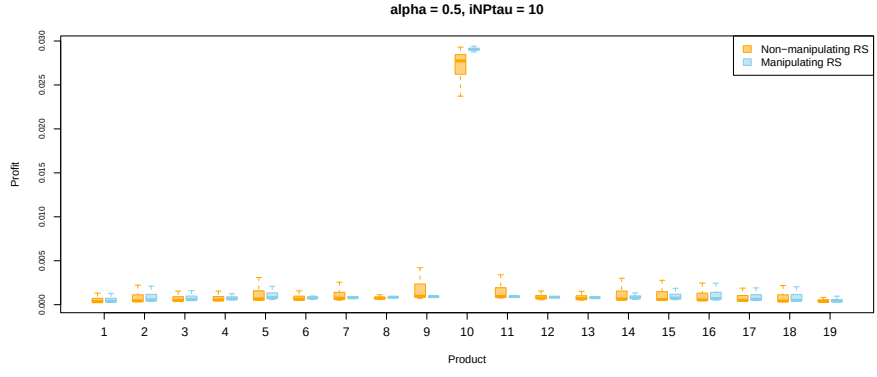

FIGURE 18.—Impact of 50% manipulation on profits for all products with product 10 favored by the recommender system ($\alpha = 0.5$).

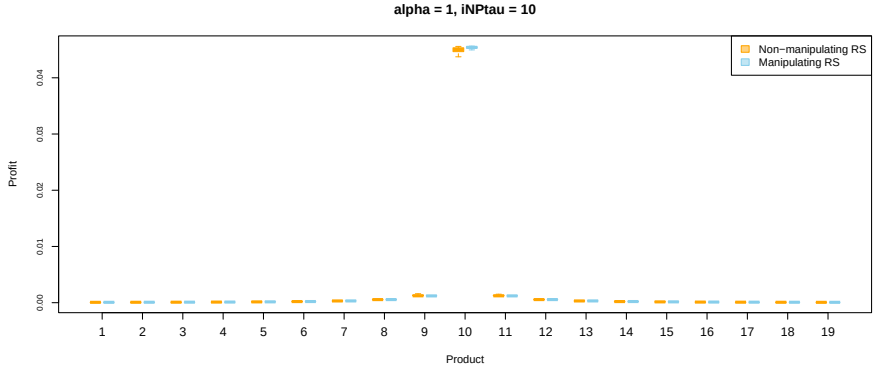

FIGURE 19.—Impact of 50% manipulation on profits for all products with product 10 favored by the recommender system ($\alpha = 1$).

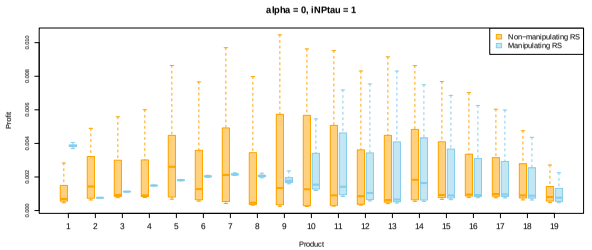

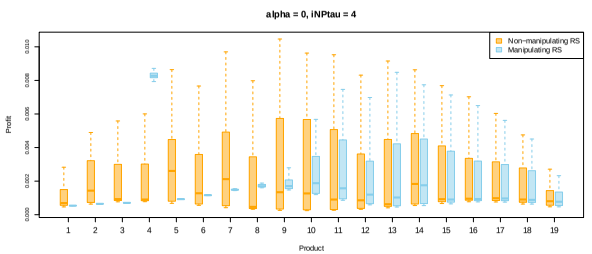

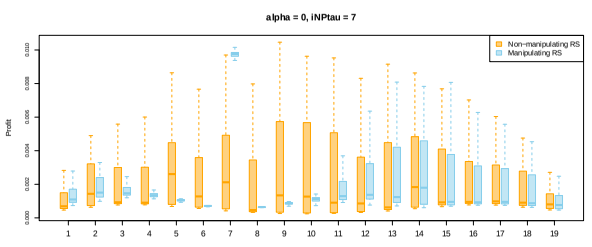

(a)

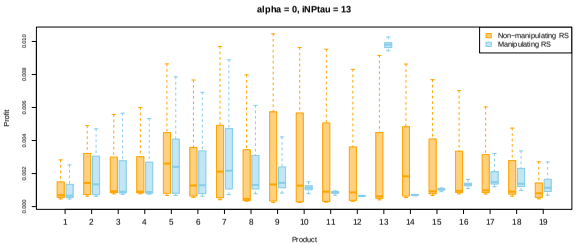

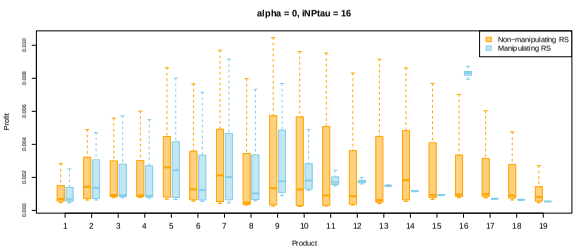

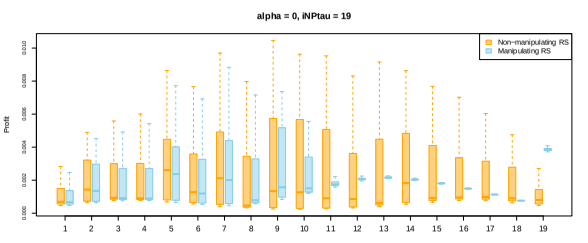

(b)

FIGURE 20.—Effect of manipulation on equilibrium profits, when the favored product differs from the central one.'INPtau' is the favored product's identity. Our labeling convention designates products based on their location, with products 1 and 19 as peripheral and product 10 as central. Key this figure: parameters $\alpha = 0$, and manipulation frequency maintained at 50%.

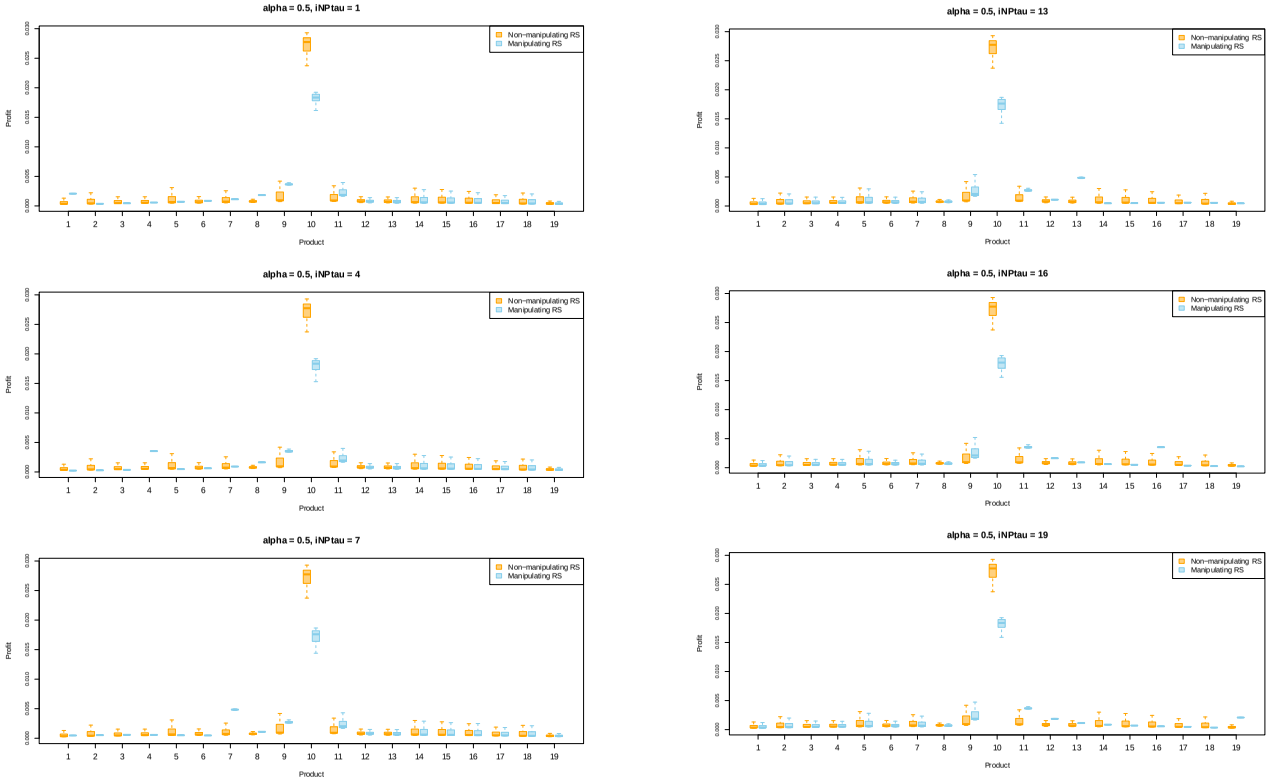

(a) (b)

FIGURE 21.—Effect of manipulation on equilibrium profits, when the favored product differs from the central one.'INPtau' is the favored product's identity. Our labeling convention designates products based on their location, with products 1 and 19 as peripheral and product 10 as central. Key this figure: parameters $\alpha = 0.5$, and manipulation frequency maintained at 50%.

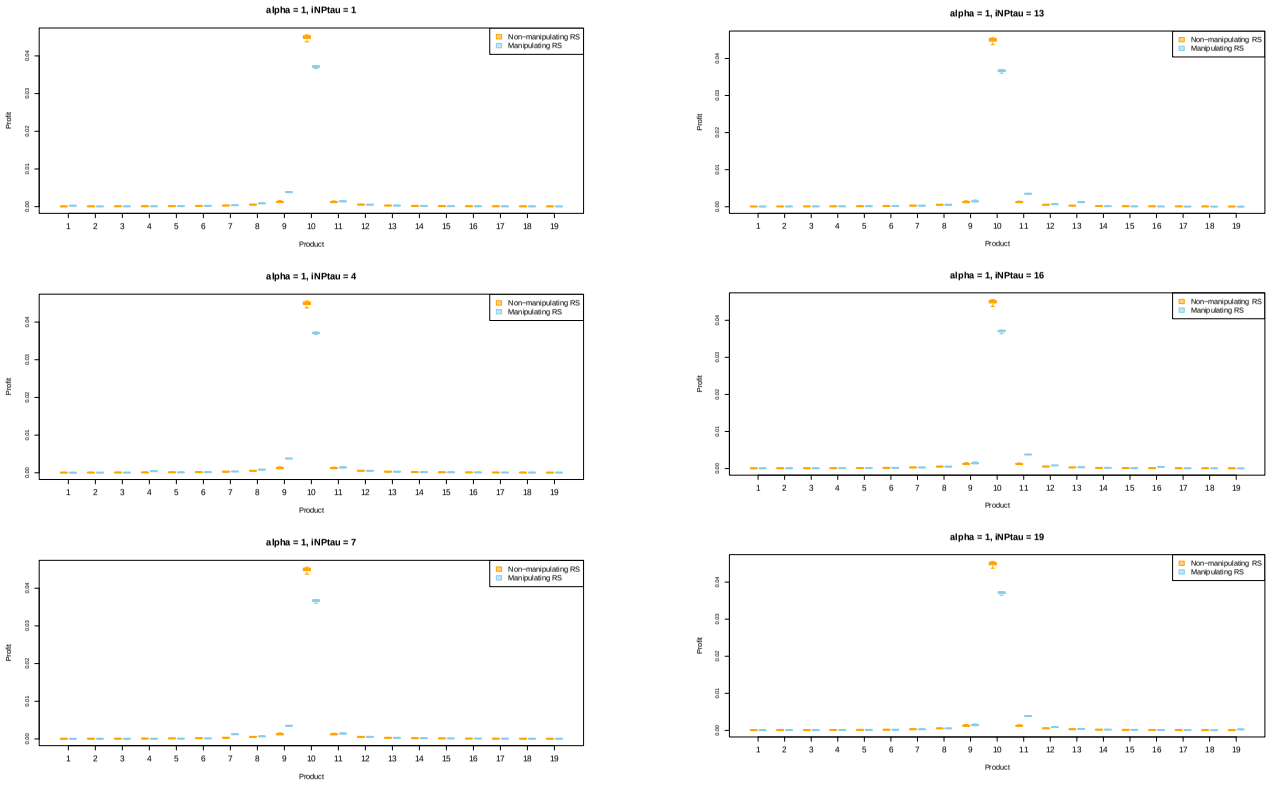

(a) (b)

FIGURE 22.—Effect of manipulation on equilibrium profits, when the favored product differs from the central one.'INPtau' is the favored product's identity. Our labeling convention designates products based on their location, with products 1 and 19 as peripheral and product 10 as central. Key this figure: parameters $\alpha = 1$, and manipulation frequency maintained at 50%.

9.3. *Consumer surplus*

In Figure 23, consumer surplus decreases when recommendations are manipulated relative to the case where recommendations are sincere. This occurs independently of which product is favored with the manipulation and the level of α.

Note that when α is large, the scope for manipulation is quite limited even with a 50% manipulation rate. This is because in this case, most consumers are already directed towards the central product, regardless of manipulation. As a result, the impact of manipulation is almost imperceptible when α is close to 1.

This change is again driven by three factors: the quality of matching, search costs, and prices. Manipulation negatively affects the quality of matching, leading to an increase in search costs for consumers. However, it also causes a substantial decrease in prices which mitigates its negative impact on consumer surplus. The scenario in which Product 10 is favored and α is set to 1 holds a distinct significance, as manipulation enhances matching, leading to improved matches and consequently, an increase in surplus. All these figures confirm the results discussed in the main text.

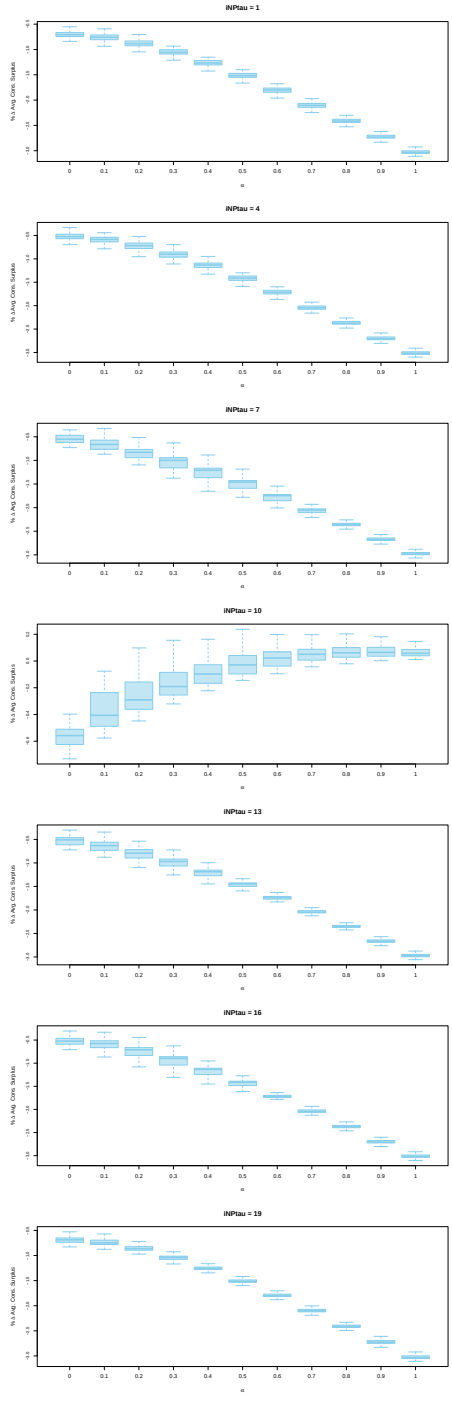

FIGURE 23.—Percentage difference in the average consumer surplus between the cases of manipulated and genuine recommendations with a rate of manipulation of 50%, for different manipulated products (labeled INPtau) and different values of α.