3-2021

The Influence of Food Recommendations: Evidence from a Randomized Field Experiment

Kamal Bookwala

Caleb Gallemore

Joaquín Gómez-Miñambres

The Influence of Food Recommendations: Evidence from a Randomized Field Experiment

The Influence of Food Recommendations:
Evidence from a Randomized Field Experiment

Kamal Bookwala, Caleb Gallemore and Joaquín Gómez-Miñambres[1]

March 2021

Abstract

We report results from a randomized field experiment conducted at two food festivals. Our primary aim is to assess the impact of two types of recommendations commonly observed in food settings: *most popular* and *chef's choice*. Subjects select a cupcake from a binary menu. The two options, offered by the same bakery, are the best seller in the bakery and the baker's recommended cupcake. Our treatments manipulate whether the recommendation is disclosed in tandem with the cupcakes in the menu. We find that the *most popular* is the only recommendation that statistically significantly increased consumers' demand relative to a baseline without recommendations. Furthermore, we find that this effect only holds for subjects from outside the local region. Our results are consistent with laboratory studies indicating information on peers' choices is a powerful influence on consumers' decisions, especially in the absence of prior knowledge.

KEYWORDS: Recommendations; Social learning; Herd behavior; Peer effects

JEL: C93, D12, D83

1. Introduction

Modern economies are characterized by an often overwhelming array of choices, which may explain why consumers rely on recommendations and endorsements when making decisions.[2] Firms, in turn, often use recommendations to influence consumers. A common strategy is to promote a product as the leading, best-selling, or most popular one on the market (Cialdini and Goldstein 2002). The returns to this "social validation" marketing strategy can be substantial. For example, appearing in the *New York Times* bestseller list generates over a 50% increase in sales for new authors (Sorensen 2007). Another approach uses "expert" recommendations from individuals perceived as having superior knowledge about a particular product. Achieving a *Michelin Star*, for example, allows restaurants to increase their prices by approximately 30% (Gergaud et al. 2015).

While several scholars have studied the marketing implications of various product recommendations (e.g., Berger, Sorensen and Rasmussen 2010; Chen and Xie 2005; Gergaud et al. 2015; Sorensen 2007), an important but underexplored area is the study of what type of recommendation is most effective in influencing which consumers' choices. For example, if a product is both popular and expertly recommended, what piece of information should a marketing campaign emphasize - and to whom? Should a "chef's choice" recommendation be included in a restaurant's menu? What types of consumers are more likely to be influenced by such recommendations? It is important to develop a better understanding of how firms can incorporate recommendations in their marketing strategies to most effectively impact consumers' choices and maximize profit.

The dearth of literature on these issues likely reflects an inherent problem in measuring recommendations' effects on consumer choices - recommendations and sales are simultaneous. This is an obvious challenge for assessing the impact of "most popular" recommendations, as recommended products are, by definition, also the most successful alternatives. A similar selection bias also affects assessments of "expert" recommendations. Products receiving such recommendations should, in principle, be of particularly high quality and hence, naturally more popular than unrecommended products, again making it difficult to distinguish recommendations' effects from the effect of the product's intrinsic characteristics. These problems become even harder if we wish to contrast how different types of information (e.g., most popular, or expert's choice) sway consumers' choices. A further challenge for assessing both types of effects is that sellers might engage in strategic signaling, displaying recommendation information only for those products that they wish to promote, which may affect consumers' beliefs about inherent quality.

To address these problems, we analyze the results of a randomized field experiment conducted at two food festivals in eastern Pennsylvania. In our experiment, fairgoers responded

---

[2] For instance, 35% of all sales at Amazon result from recommendations, while 75% of the content watched in Netflix has been suggested by its recommendation system (Amatriain and Basilico 2012).

to a survey about the festival and, as a reward for participating, were offered a choice between two cupcakes from a local bakery. One option was the bakery's best seller (a "most popular" condition), while the other option was the baker's favorite cupcake (a chef's choice, or in this case, "baker's choice" condition). In two of our treatments, we varied the display of each cupcake's recommendation type jointly with a "thumbs up" symbol. Since it is possible that a recommendation's effect is due simply to the salience it grants the accompanying menu item, we also included two additional treatments in which only the thumbs-up symbol accompanied the cupcakes. This allowed us to compare differences in subjects' choices resulting from informative, as opposed to ambiguous, recommendations. Moreover, since the cupcakes were offered as a reward, our experimental design controlled for the possible strategic use that the seller is signaling the products that she is trying to sell.

Our results indicate that only the "most popular" recommendation statistically significantly affected consumer's choices, its selection rate increasing 13% under the explicit recommendation condition. Neither the ambiguous recommendations (thumbs-up alone), nor the "baker's choice" significantly influenced consumers' choices. We also find some evidence that being located first in the menu increases consumers' demand, but this order effect disappears when the product recommendation is informative, rather than ambiguous. Therefore, informative recommendations may mitigate order effects in menus. Finally, we find that both the informative recommendation effect and the order effect hold only for subjects whose residences were in zip codes outside the local region, who were more likely to lack prior knowledge about the products. This shows that subjects' behavior is likely affected by the novel information content of the recommendation rather than the result of a conformity effect.[3]

2. Related Literature

Our study relates to work investigating the influence of: (i) peers, and (ii) experts on consumer's behavior. In the first category, we have studies on what psychologists first referred to as "social learning" (Bandura 1977).[4] In economics, the seminal papers on this topic are Banerjee

---

[3] A *conformity effect* exists when individuals adopt the observed choices of others, not because of the information it conveys, but because they want to conform (Asch, 1955; Cialdini and Goldstein, 2004; Bernheim 1994). Some recent experiments in the economics literature have provided clear evidence of preference conformism (e.g., Fatas, Hargreaves-Heap and Rojo-Arjona, 2018). In our experiment, it is indeed possible that a fraction of subjects are conformists and, when told that other people selected a cupcake, they derive utility from choosing the same. Preference conformism is consistent with our key result that only the most popular recommendation is effective. However, our finding that only non-locals, who are more likely to lack prior knowledge about the alternatives, are affected by the most popular recommendation suggests that subjects' choices are also influenced by the information content of the recommendation. In short, while we cannot rule out the possibility that some subjects are conformists, our results indicate that local subjects' choices are not affected by recommendations, possibly because they have (or think they have) superior knowledge about the products.

[4] There are two types of "social learning". Informational learning refers to situations where individuals communicate and share information in a personal manner. Under observational learning, on the other hand, behavior is influenced by the mere observation of one's choices (Cai et al. 2009). Because, in our experiment, subjects make decisions individually and without communication, the former cannot occur and hence our study focuses on observational learning.

3

(1992) and Bikhchandani, Hirshleifer, and Welch (1992), who examine how information transmission among consumers may encourage herding behavior, prompting individuals to preferentially choose what others have chosen, even if they have a private signal that they would like another product better.[5] Evidence of social learning was first documented in the lab (e.g., Anderson and Holt 1997), while some field studies have appeared more recently. For example, Salganik, Dodds and Watts (2006) created a music market using a database of unknown songs and artists and manipulated whether the participants saw the number of downloads made by others on a particular song. Learning this information made popular songs more popular and unpopular songs less popular, increasing the inequality of market outcomes. Similarly, Moretti (2011) found that movies whose box-office sales during opening weekend were higher than expected performed better in future weeks, and Anderson and Magruder (2012) showed that positive consumer evaluations of restaurants on Yelp.com induced more table reservations. Peer effects have also been reported in other domains such as the diffusion of residential solar panels (Bollinger and Gillingham, 2012), drinking of alcoholic beverages (Deconinck and Swinnen 2015), weight-related behavior (Ali, Amialchuk, Heiland 2011) and financial decisions (Bursztyn, Ederer, Ferman, and Yuchtman 2014).

Perhaps the most closely related work in this stream of literature is Cai, Chen, and Fang (2009). Using a randomized field experiment, they studied how differing restaurant menu features affected diners' choices. When provided with a list of the previous week's "top selling dishes," consumers were more likely to select those options, but the effect disappeared if the same items were listed merely as "sample dishes". To our knowledge Cai et al. (2009) is the only paper besides ours attempting to assess the degree to which the "saliency effect" accounts for recommendations' impacts. Nevertheless, our work differs from theirs in several respects. First, we compare the effect of different types of recommendations on demand. Moreover, the way we control for saliency is quite different because our "thumbs-up" recommendations signal that the product was still recommended but without indicating the recommendation's source. Therefore, our design allows us to compare recommendations' saliency effects to effects arising from the type of information the recommendation conveys. Finally, as opposed to Cai et al. (2009), who rely on a two-stage randomization strategy (first restaurant, then table), we use a cleaner randomization at the individual level, and likely have access to a more diverse population as a result of conducting the experiment at two outdoor food festivals.

Our study also relates to work examining how *expert opinion* affects consumers' decisions in the face of information asymmetries. Results from this literature show that expert recommendations increase demand for products such as wine (Friberg and Grönqvist 2012; Hilger, Rafert and Villas-Boas 2019), fine dining (Gergaud, Storchmann and Verardi 2015), books (Sorensen 2007; Berger, Sorensen, and Rasmussen 2010), and movies (Ginsburgh 2003;

---

[5] An alternative explanation for why individual demand is affected by peers' action is network externalities. For example, in Becker (1991) the demand for a good (e.g., a restaurant meal) depends positively on its aggregate quantity demanded because people enjoy consumption more when others are consuming as well.

Reinstein and Snyder 2005). Psychologists have also argued that authority figures labeled as "experts" might affect how subjects evaluate recommendations (e.g., Cialdini and Goldstein 2004), with impartial experts being more effective than those with a stake in consumers' choices (Cialdini and Goldstein 2002). Crucially, unlike our controlled field setting, most existing literature on expert recommendations relies on quasi-experimental empirical approaches to avoid the problem of spurious correlation between recommendations and demand. One exception is Hilger et al. (2011), who uses a randomized field experiment to show that positive expert reviews increase the demand for wine. Unlike our study, however, they do not consider the effect of different types of recommendations or compare explicit and ambiguous recommendations.

3. The Experiment

Our study was part of a larger marketing project conducted in conjunction with Lafayette College, the Greater Easton Development Partnership, and the Lehigh Valley Research Council, an academia-government research agency connecting local researchers to nonprofit and local government agencies. The project consisted of an in-person survey conducted with attendees of two local food festivals: Garlic Fest (10/5-10/6/19) and Bacon Fest (11/2-11/3/19). At the end of the survey, and as a reward for participating, participants were offered a bottle of water and a cupcake from a local bakery. They made their choice of cupcake by circling one of two options from a paper menu and they received the cupcake of their choice immediately afterwards. Each menu contained the names, descriptions, and pictures of both cupcakes (see Figure 1). We used this part of the survey to conduct our randomized field experiment.
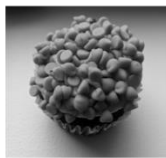

Figure 1. Baseline menu. Sweet Girlz Bakery logo used with permission of the owner.

The two cupcake flavors on offer were "Peanut Butter Bomb" (PBB) and "French Toast" (FT). While the bakery reported that both were relatively popular cupcakes, PBB was the most popular, and FT was the baker's favorite.[6] At the time the experiment was conducted the bakery did not use any recommendation system of their own. Both cupcakes had the same retail price ($2.50) as well as similar size and calorie content. Since the menus were attached to the surveys, we were able to match cupcake choices with some demographic information about the participants, allowing us to control for differences in cupcake preferences across social groups.

3.1.*Experimental treatments*

Our experiment followed a between-subjects design consisting of a control with no recommendations (see Figure 1) and four recommendation treatments: *baker's choice*; *most popular*; *thumbs-up FT* and *thumbs-up PBB*. The first two treatments (baker's choice and most popular) manipulated whether each recommendation type appeared next to the corresponding cupcake or not. The information was accompanied by a thumbs-up symbol. In the last two treatments only the thumbs-up symbol appeared. Therefore, the recommendation in these treatments was left deliberately ambiguous, so subjects did not know why the cupcake was recommended. The purpose of the 'thumbs-up' treatments was to test whether the effectiveness of the recommendation came from the information it conveys (baker's choice/most popular) or simply because a recommendation of any kind increases the option's salience. Finally, to control for possible order effects we randomized the order in which the cupcakes appeared in the menus for all treatments. In total, this amounted to 10 different experimental menu conditions.

Figure 2: Experimental Treatments
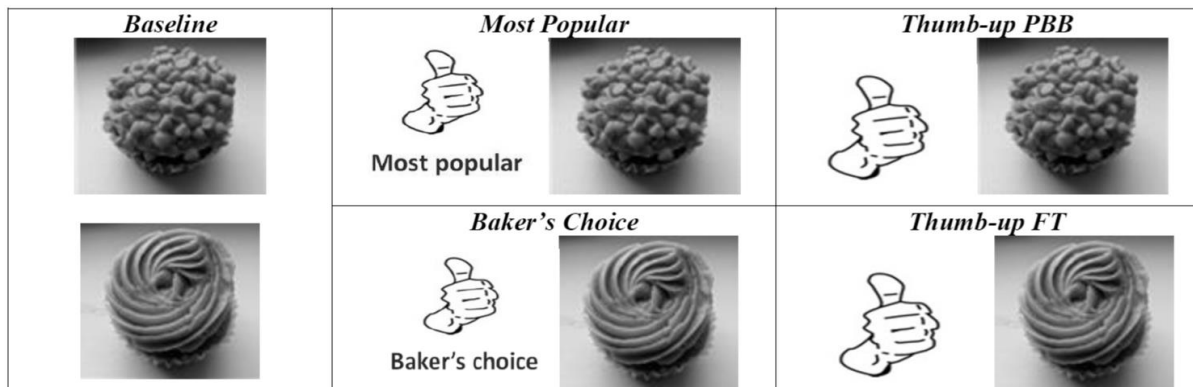
Note: Each menu consisted of pictures of the two cupcakes and their descriptions (see Figure 1 for Baseline). The treatments manipulated whether each recommendation type appeared next to the corresponding cupcake or not. We only displayed one recommendation per treatment. We show all the recommendation treatments in Figure A1 and Figure A2 in the appendix.

---

[6] We gathered this information in a face-to-face informal conversation with the owner, the baker, and other workers at Sweet Girlz Bakery before explaining the purpose of the study (or even that we were conducting a study at all). Therefore, it is unlikely that their answers were motivated by a desire to advertise a particular cupcake.

3.2. *Procedures and randomization*

The two food festivals where we conducted the study took place in Easton, PA, on two different weekends. We collected data on both days of each festival (Saturday and Sunday) between 11am and 4pm. Both festivals were held around the city's main square and its four adjacent streets. We hired eight college students for each data collection day and divided them into four teams of two to distribute surveys. Each team was stationed on one of the four streets that feed into the city circle (see Figure A3 in the appendix). The survey teams were equipped with a utility wagon containing surveys, informed consent forms, pens, clipboards, water bottles, and, of course, cupcakes (see Figure A4 in the appendix). We shuffled all menu conditions in the pile of surveys we gave to each team such that all conditions were equally likely to be distributed in all locations at all times. Survey organizers were available to the interviewers during data collection at a home base location in the square and checked in with the interviewing teams approximately once every half hour to answer questions or replenish any survey supplies that were running low.

A week before the festivals, the interviewers had an orientation meeting with the survey organizers covering project logistics. During data collection, interviewers approached festival-goers and asked them if they would be interested in participating in a brief survey and getting a cupcake as a reward. The interviewers let participants answer the questions alone, instructing each participant to take the survey individually but providing explanation and clarification as necessary.[7] When finished, the participants handed the survey back to the interviewers and received a water bottle and the cupcake that they had selected from the menu.

We collected a total of 859 surveys over the course of both food festivals. Omitting 28 subjects who did not select a cupcake from the menu left 831 usable observations. The demographic composition in both festivals was similar. Most of the subjects were female (≈ 60%), under the age of 34 (≈ 50%), Caucasian (≈75%), and had a household income above $75,000 (≈ 50%). A slight majority of the subjects were from the local area, which we defined as zip codes in Pennsylvania's Lehigh Valley and along New Jersey's I-78 corridor from the state line to Clinton, New Jersey (≈ 56.7%). We show participants' demographics in more detail in Figure 3.

---

[7] Despite the instruction of answering the survey alone, we cannot control whether some participants who attended the festival in groups did not engage in communication with their partners or chose the cupcake for somebody else. Our randomization ensures that this confounding effect is equally present in all treatments. However, it is likely that our results would have been stronger if all subjects made decisions in complete isolation.
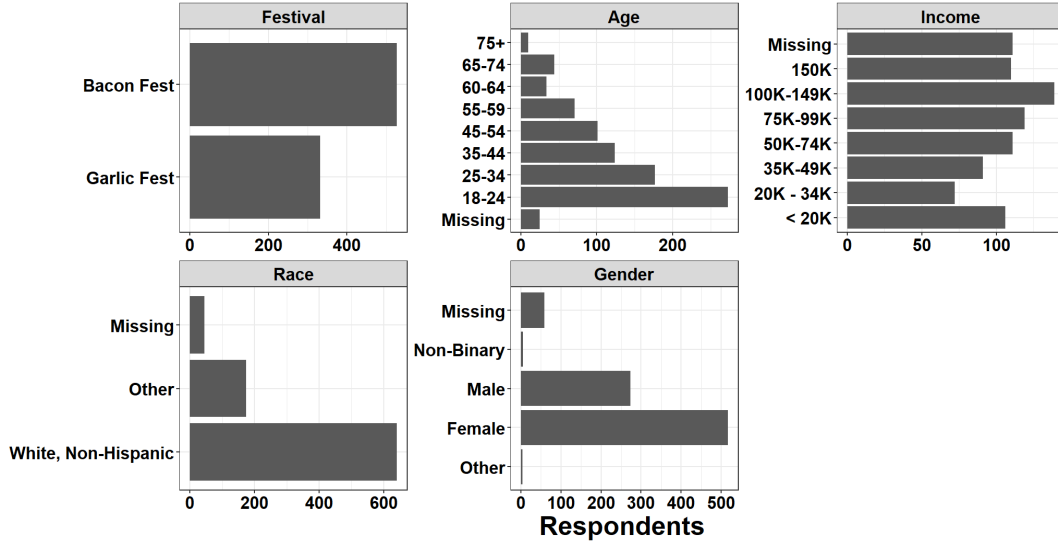
Figure 3. Subjects' demographics.

4. Results

4.1. The effectiveness of the recommendation: most popular vs. baker's choice

The main goal of this study is to investigate the effectiveness of the "most popular" and "baker's choice" recommendations. As we can see in Table 1, the selection rate of the recommended cupcake (the PBB) in the most popular treatment was 57%, a statistically and substantively significant increase relative to the 45% selection rate in the baseline (two-sided t test, $p=0.0286$). However, none of the other treatments affected participants' choices significantly. As it so happened, the PBB, the most popular cupcake in the bakery, was not the cupcake most often selected in our baseline condition. It became the most popular choice only when subjects were told it was the most popular. In other words, the most popular recommendation created a self-fulfilling prophecy.[8]

Table 1: Selection rate of cupcakes across treatments

	Baseline (N=153)	Thumbs-up FT (N=166)	Thumbs-up PBB (N=166)	Baker's choice (N=160)	Most Popular (N=184)
PBB	45%	45%	50%	48%	57%
FT	55%	55%	50%	53%	43%
Std. Error	4.0%	3.9%	3.9%	3.9%	3.6%

While this preliminary analysis is intriguing, it is necessary to control for several possible confounders, particularly order effects and differences in preferences across social groups, to be

---

[8] Salganik and Watts (2008) report a self-fulling property of popularity information in an artificial music market experiment that, unlike our study, relied on subjects' deception. After inverting the true popularity of a list of songs, the authors found that even if the information about popularity is initially wrong, it became real over time.

confident these differences are not simply random chance. To control for potential confounders, we estimated two logistic regression models, presented in Table 2. The dependent variable takes the value one if the subject chose the PBB cupcake. These models confirm the initial finding that the most popular recommendation condition was the only one that statistically significantly affected the frequency of consumers' choices relative to the baseline. This result holds when we control for participants' demographics (Model 2), and the estimated effect actually becomes slightly stronger. We also find evidence of order effects - subjects are more likely to choose the cupcake listed first. This pattern also remains when we include the demographic controls.

Table 2: Logistic regressions for PBB choice.

Dependent variable: *PBB*	Model 1: Basic Model	Model 2: Model with Demographic Controls
(Intercept)	-0.248 (0.178)	-0.650* (0.258)
Thumbs Up FT	0.210 (0.226)	0.237 (0.239)
Thumbs Up PBB	-0.0208 (0.226)	-0.0577 (0.242)
Baker's Choice	0.0986 (0.228)	0.258 (0.243)
Most Popular	0.516* (0.222)	0.604* (0.236)
PBB Listed Second	-0.331* (0.141)	-0.332* (0.150)
White, Non-Hispanic		0.576** (0.187)
Male		0.0591 (0.160)
Income Above $100K		0.0756 (0.164)
Surveyed at Garlic Fest		0.0816 (0.154)
Area Under the ROC Curve	0.573	0.602
N	825	747

Note: Coefficients on log-odds scale. Standard errors in parentheses.

***p<.001, **p<.01, *p<.05

To make the substantive effects of the experimental conditions clear, we compute the predicted probability that a subject selects PBB in the different experimental conditions, with different item orders, based on the coefficients in Model 2. We present these estimates in Figure 4.
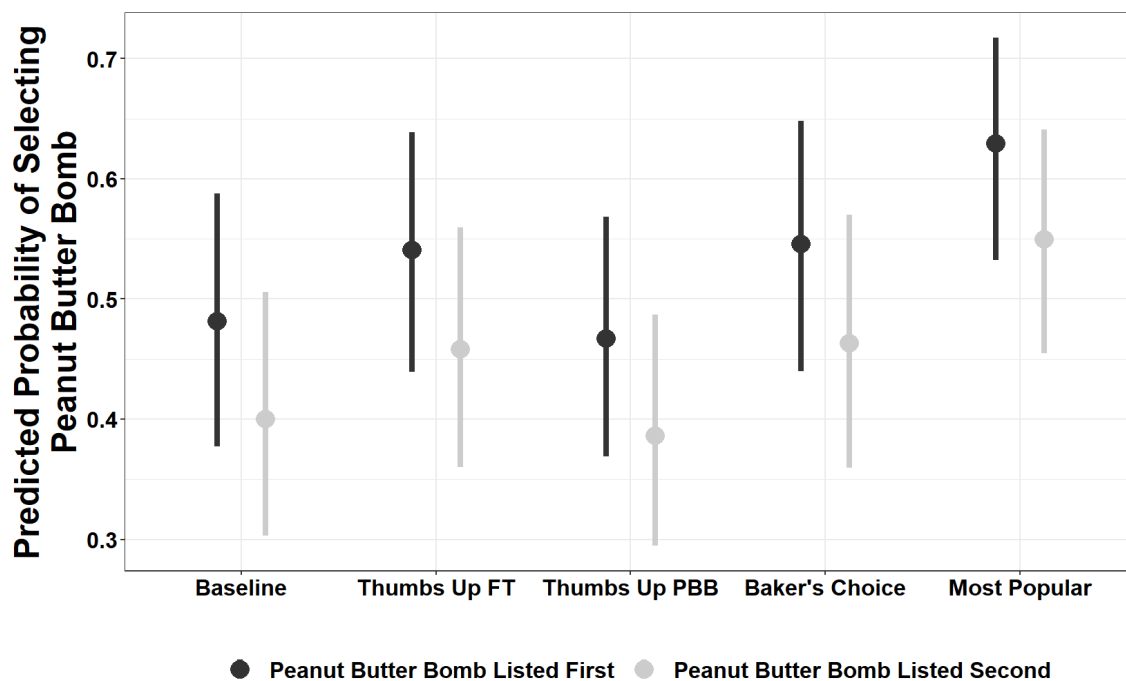
Figure 4: Predicted probability, with 95% confidence intervals, of selecting the PBB cupcake. Based on Model 2 coefficient estimates. All other variables set to their mode.

As Figure 4 shows, the "*most popular*" condition generates a substantively significant difference in the probability of selecting PBB (the recommended cupcake) relative to the baseline condition. Respondents in this condition are predicted to be about 13 percentage points more likely to select PBB than those in the baseline. The predicted probability of selecting PBB in the condition with only the thumbs up sign, by contrast, is virtually identical to the baseline probability. This evidence indicates that the effectiveness of the "most popular" recommendation comes from the information it conveys and not from a cupcake-specific effect of the recommendation.

An additional interesting finding is that the order effect appears to be stronger when considering ambiguous versus informative recommendations (see Table 3). While the coefficient for PBB being listed second on the menu remains negative when considering only the baseline and informative recommendation conditions, it is no longer statistically significantly different from zero, and the estimated order effect is substantively smaller. This result could indicate that, while ambiguous recommendations do not provide sufficient information to overcome simple decision heuristics, informative recommendations can overcome order effects to some extent.[9]

---

[9] While the saliency of being listed first should be stronger in a menu with more than two alternatives; in our environment, order effects might be playing a role in the baseline because being listed first can be interpreted as an implicit signal of being recommended. This would explain why, when clear recommendations are present, order effects disappear.

Table 3: Logistic regressions for PBB choice, comparing models with informative versus ambiguous recommendations.

Dependent variable: *PBB*	Model 3: Informative Recommendations	Model 4: Ambiguous Recommendations
(Intercept)	-0.538 (0.303)	-0.896** (0.310)
Thumbs Up FT		0.264 (0.245)
Thumbs Up PBB		-0.0522 (0.245)
Baker's Choice	0.242 (0.243)	
Most Popular	0.595* (0.236)	
PBB Listed Second	-0.327 (0.195)	-0.401* (0.198)
White, Non-Hispanic	0.424 (0.239)	0.846*** (0.257)
Male	0.327 (0.204)	-0.0668 (0.216)
Income Above $100K	0.0172 (0.210)	0.282 (0.217)
Surveyed at Garlic Fest	-0.0762 (0.199)	0.150 (0.205)
Area Under the ROC Curve	0.610	0.625
N	448	439

Note: Coefficients on log-odds scale. Standard errors in parentheses.

***p<.001, **p<.01, *p<.05

Result 1 (*most popular vs. baker's choice*)

i) The most popular recommendation significantly increased the cupcake's choice rate. Neither the chef's choice recommendation nor the thumbs-up symbol alone affected subjects behavior between treatments.

ii) A cupcake was more likely to be chosen when listed first, but this order effect is only statistically significant in the thumbs-up alone treatments (ambiguous recommendations).

4.2. Local vs. non-local subjects

To finish, we consider whether recommendations have a different effect for subjects who reside outside of the region where the experiment took place. Since local subjects are more likely to be familiar with the bakery that provided the cupcakes (a small local business) than those from outside the region, this analysis allows us to test whether prior knowledge about the products alters recommendations' effectiveness. For this analysis, we used two different techniques to divide our sample into two groups. The first approach drew on our own local knowledge of the areas generally perceived to be part of the region. For this approach, we defined the local region to include all zip codes in Pennsylvania's Lehigh Valley and along New Jersey's I-78 corridor from the Pennsylvania-New Jersey state line to Clinton, New Jersey (N = 487). The other group included all subjects reporting a residence outside this area (N = 291).[10] Because this approach relies on a subjective definition of region, we also took an objective approach, using the Google Maps API to estimate the shortest driving time between the geographic center of each subjects' zip code and the main parking structure in downtown Easton. We then split the sample into respondents living in zip codes below the median drive time in the sample and those at or above the median, with those at or above the median travel time deemed to be outside the region. In Figure 5, we show the geographic distribution of respondents according to their residence.
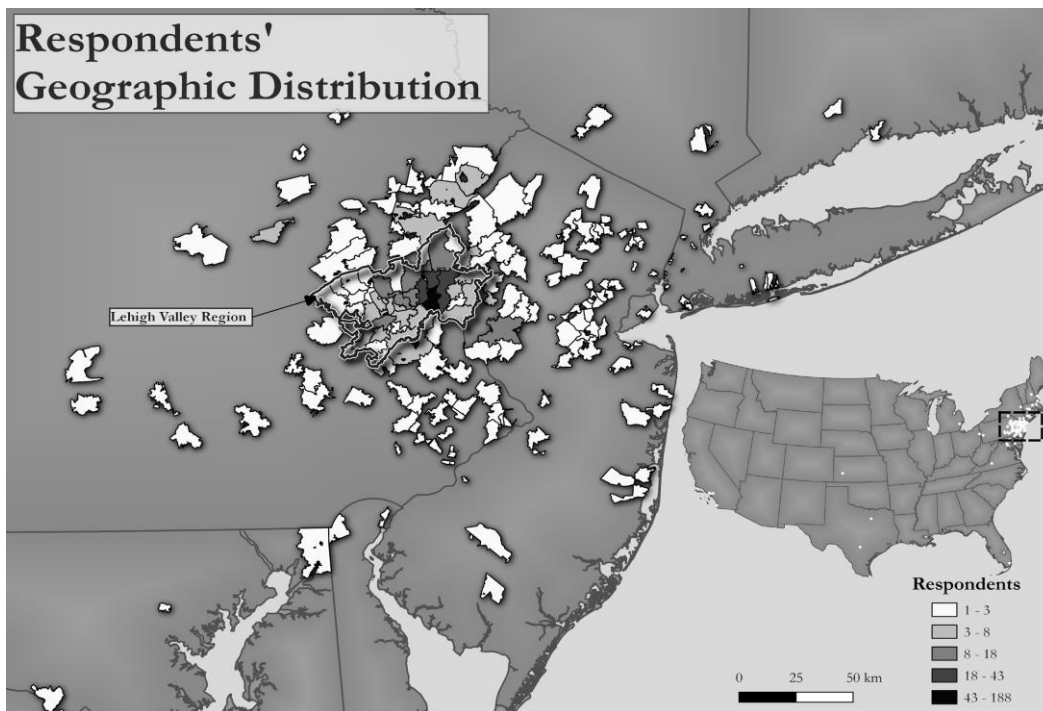
Figure 5. Subjects' Geographic Distribution, by zip code.
Note: Projection for main map: State Plane Pennsylvania South; Projection for inset: US National Atlas Equal Area.

In Table 4 below we show the main results of this analysis.

---

[10] Our results are robust to different specifications of region, such as residence within versus outside Easton, PA, city limits.

Table 4: Logistic regressions for PBB choice, interacting experimental condition with residence inside or outside the region.

Dependent variable: *PBB*	Model 5: Outside Region Defined as Outside Lehigh Valley and I-78 Corridor in New Jersey	Model 6: Outside Region Defined as at or Above Median Drive Time to Downtown Easton
(Intercept)	-0.399 (0.287)	-0.297 (0.234)
Thumbs Up FT	0.0791 (0.297)	0.0777 (0.363)
Thumbs Up PBB	-0.295 (0.296)	-0.517 (0.368)
Baker's Choice	-0.0835 (0.299)	-0.126 (0.355)
Most Popular	0.201 (0.289)	-0.00914 (0.355)
Thumbs Up FT * Outside Region	0.502 (0.511)	0.314 (0.516)
Thumbs Up PBB * Outside Region	0.725 (0.519)	0.782 (0.524)
Baker's Choice * Outside Region	1.02 (0.525)	0.586 (0.522)
Most Popular * Outside Region	1.21* (0.510)	1.13* (0.509)
Outside Region	-0.702 (0.380)	-0.695 (0.379)
PBB Listed Second	-0.341* (0.151)	-0.315* (0.161)
White, Non-Hispanic	0.546** (0.189)	0.534** (0.200)
Male	0.0947 (0.162)	0.107 (0.173)
Income Above $100K	0.0551 (0.165)	0.221 (0.176)
Surveyed at Garlic Fest	0.0925 (0.156)	0.0949 (0.167)
Area Under the ROC Curve	0.617	0.616
N	747	663

Interestingly, we find that the "Most Popular" condition has a much more pronounced effect for non-locals, regardless of the method used to define locality, than for the sample as a whole. Indeed, when we interact the experimental conditions with the local/non-local categories, the coefficient for the "Most Popular" condition is no longer statistically significant, indicating there is no detectable effect for regional residents. In other words, respondents from outside the region, however defined, are the primary drivers of the observed effect of the "Most Popular" experimental condition. Figure 5 presents predicted probability estimates for these models for comparison. It is particularly notable that the magnitude of the "Most Popular" effect is substantially stronger for non-locals than for the population as a whole. As can be seen in Figure 5, respondents from outside the Lehigh Valley and the I-78 corridor in the "Most Popular" condition were roughly 30% more likely to choose PBB than in their baseline condition, while those at or above the median drive time were roughly 25% more likely to choose PBB than in their baseline. This finding strongly suggests that "Most Popular" messaging can be very effective with audiences lacking relevant information.

Surprisingly, it also seems that the "Baker's Choice" is associated with an increase in the odds that subjects selected PBB, though this effect is marginally outside of our confidence interval (p-value=0.052). One possible reason for this marginally significant negative effect of the "Baker's Choice" condition might be the subjects' perception of a conflict of interests in this recommendation. While our experimental design (where cupcakes were offered as a reward for completing the survey rather than sold) should control for the possible strategic use that the baker is signalling the product she is trying to sell, it is possible that some subjects believed the baker had something to gain by influencing them to choose a particular cupcake. In fact, Cialdini and Goldstein (2002) explicitly mention the "chef's choice" as an example of a recommendation that it is likely to be ineffective because of this conflict of interest. Similarly, findings in the "cheap talk" literature indicate that decision makers anticipate the possible bias in experts' recommendations (Chakraborty and Harbaugh 2010). This could help explain why non-locals, who might be expected to have less social trust in the local business, seem to have reacted against the "Baker's Choice" recommendation.
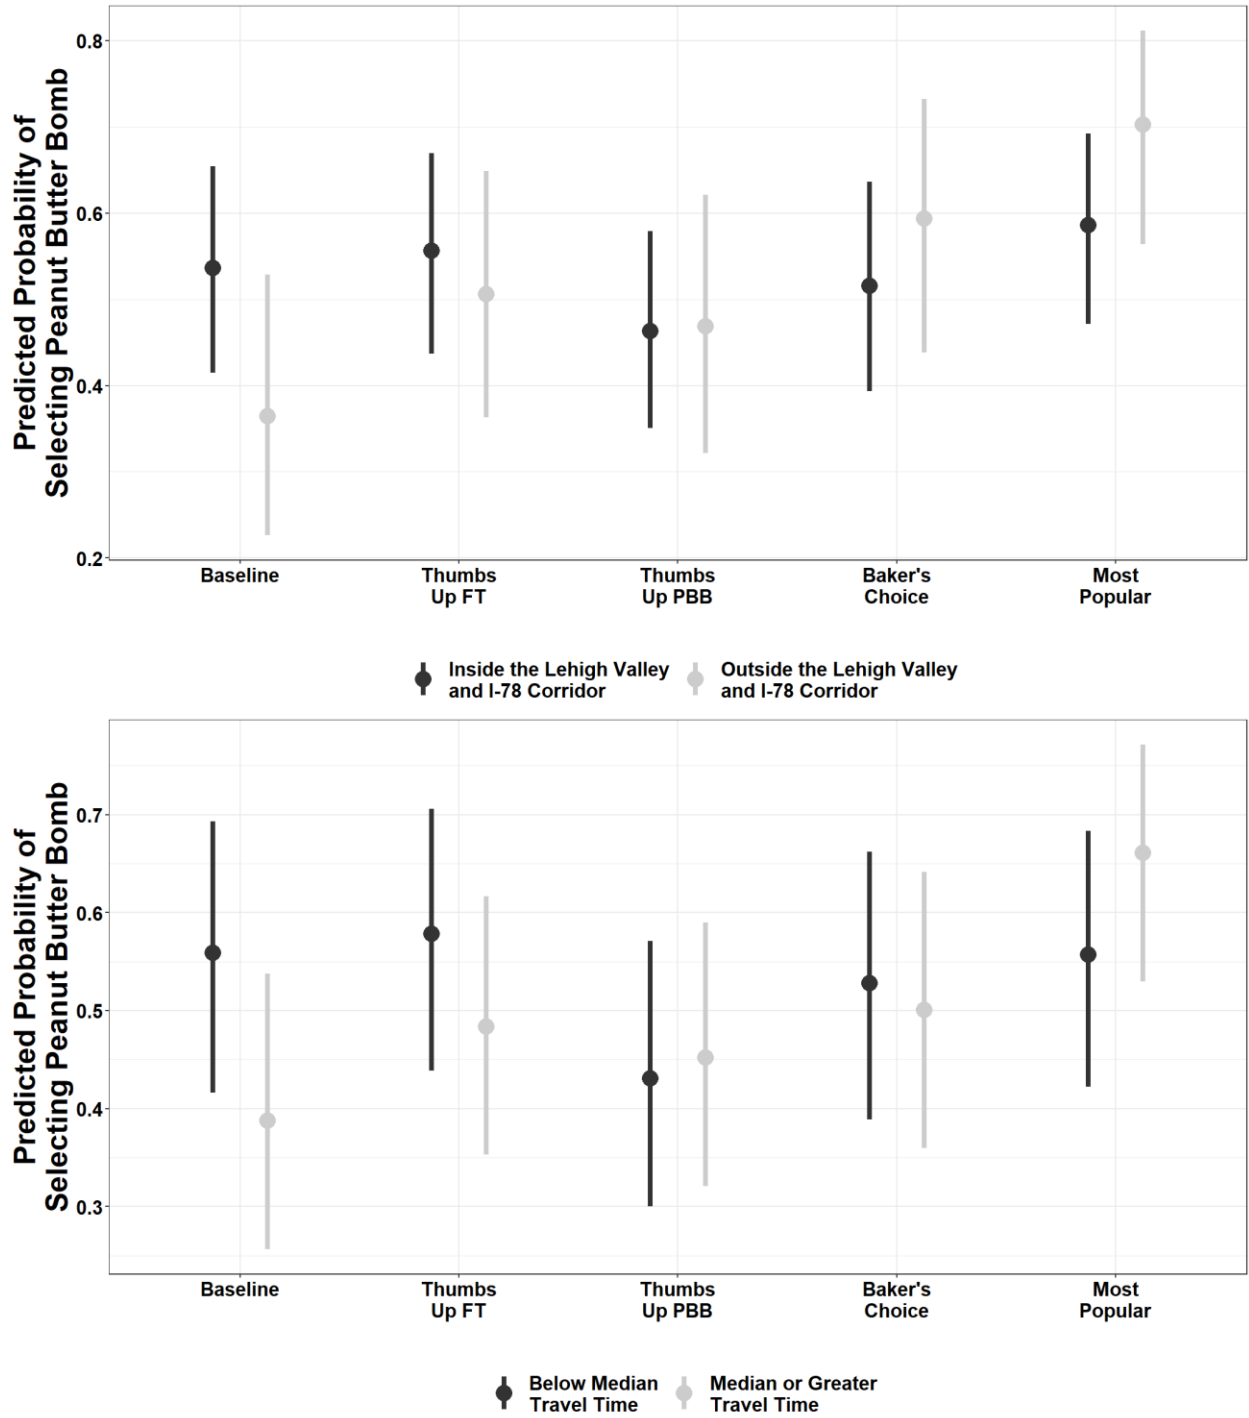
Figure 5: Predicted probability of selecting PBB, with 95% confidence intervals, by residence zip code and estimated travel time. Based on coefficient estimates from Model 5 (top) and Model 6 (bottom).

Result 2 *(Local vs. non-local subjects).*

(i) Local subjects were not influenced by any recommendation type.

(ii) The most popular recommendation was only effective among non-locals. The magnitude of this effect was twice as strong among non-locals than in the overall population.

(ii) The baker's choice recommendation had a marginally significant negative effect on the choice rate of non-local subjects.

5. Conclusion

In this paper, we have presented the results of a randomized field experiment testing the effect of different types of food recommendations on the demand of cupcakes. We found that the "most popular" recommendation was effective in increasing the selection rate of the recommended product. However, neither a generalized, uninformative recommendation (thumbs-up symbol alone) nor the "baker's choice" conditions significantly altered consumers' choices. While the estimated impact on consumers' choices is relatively small, at an increase of around 13% in the probability of selecting the recommended cupcake, it is important to remember that this increase in certainty can come at virtually zero cost, helping clarify producers' expected sales, or, perhaps more deviously, creating beneficial self-fulfilling prophecies. Furthermore, we find a much more substantial increase of around 25-30% for non-local subjects, depending on the definition of locality, who were more likely to lack prior knowledge about the products.

Consistent with a range of studies focusing on peer effects and social learning, our findings suggest that "most popular" recommendations are an effective way of boosting demand, especially for those consumers who lack prior information about the product (such as non-locals). Nevertheless, our results also indicate that uninformative recommendations lacking a specific reason why the product is recommended (such as a thumb up symbol) are likely to be ineffective. Finally, recommendations of the "chef's choice" type might backfire and actually decrease demand slightly, possibly because of lack of trust in the seller's motives.

A better understanding of which recommendation type is most effective for different consumers can inform firms' marketing strategies involving food recommendations. In particular, our results indicate that, whenever possible, firms should prioritize favorable popularity information when designing their recommendations systems. Importantly, our finding that only non-locals are influenced by "most popular" recommendations could raise a particularly helpful strategy for smaller firms in tourism-heavy regions. Firms might be able to use this strategy to better market to potential new customers, who might be expected to be unfamiliar with menu offerings. Moreover, we have argued that recommendations from experts that might be perceived to have a stake in the product's success (e.g., "chef's choice") are likely to generate no effect at best, and perhaps even a negative response from some consumers (non-locals in our study).

This experiment suggests several promising lines for future research. In our study, we have focused on two recommendations that are commonly present in food settings: "most popular"

and "chef's choice". However, there are many other recommendation types that deserve further exploration (e.g., "latest product", "local product", "consumers who bought x also purchased y"). Moreover, the products in our study were provided to subjects free of charge, but future research might consider the interaction between prices and product recommendations. Finally, the literature on recommendations would benefit for a better understanding of how consumers perceive and interpret different product recommendations and their long-term effects in terms of consumers' satisfaction.

References

Ali, M. M., Amialchuk, A., & Heiland, F. W. (2011). Weight-related behavior among adolescents: the role of peer effects. *PloS one*, 6(6), e21179.

Amatriain, X., & Basilico, J. (2012). Netflix recommendations: Beyond the 5 stars (part 1)–the netflix tech blog. URL: http://techblog.netflix.com/2012/04/netflix-recommendations-beyond-5-stars.html.

Anderson, L. R., & Holt, C. A. (1997). Information cascades in the laboratory. *The American Economic Review*, 847-862.

Anderson, M., & Magruder, J. (2012). Learning from the crowd: Regression discontinuity estimates of the effects of an online review database. *The Economic Journal*, 122(563), 957-989.

Asch, S. E. (1955). Opinions and social pressure. *Scientific American*, 193(5), 31-35.

Bandura, A. 1977. Social Learning Theory. Englewood Cliffs, NJ: Prentice Hall.

Banerjee, A. V. (1992). A simple model of herd behavior. *The Quarterly Journal of Economics*, 107(3), 797-817.

Becker, G. S. (1991). A note on restaurant pricing and other examples of social influences on price. *Journal of Political Economy*, 99(5), 1109-1116.

Berger, J., Sorensen, A. T., & Rasmussen, S. J. (2010). Positive effects of negative publicity: When negative reviews increase sales. *Marketing Science*, 29(5), 815-827.

Bernheim, B. D. (1994). A theory of conformity. *Journal of Political Economy*, 102(5), 841-877.

Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5), 992-1026.

Bollinger, B., & Gillingham, K. (2012). Peer effects in the diffusion of solar photovoltaic panels. *Marketing Science*, 31(6), 900-912.

Bursztyn, L., Ederer, F., Ferman, B., & Yuchtman, N. (2014). Understanding mechanisms underlying peer effects: Evidence from a field experiment on financial decisions. *Econometrica*, 82(4), 1273-1301.

Cai, H., Chen, Y., & Fang, H. (2009). Observational learning: Evidence from a randomized natural field experiment. *American Economic Review*, 99(3), 864-82.

Chakraborty, A., & Harbaugh, R. (2010). Persuasion by cheap talk. *American Economic Review*, 100(5), 2361-82.

Chen, Y., & Xie, J. (2005). Third-party product review and firm marketing strategy. *Marketing Science*, 24(2), 218-240.

Cialdini, R. B., & Goldstein, N. J. (2002). The science and practice of persuasion. *Cornell Hotel and Restaurant Administration Quarterly*, 43(2), 40-50.

Cialdini, R. B., & Goldstein, N. J. (2004). Social influence: Compliance and conformity. *Annual Review of Psychology*, 55, 591-621.

Deconinck, K., & Swinnen, J. (2015). Peer effects and the rise of beer in Russia. *Food Policy*, 51, 83-96.

Fatas, E., Heap, S. P. H., & Arjona, D. R. (2018). Preference conformism: An experiment. *European Economic Review*, 105, 71-82.

Friberg, R., & Grönqvist, E. (2012). Do expert reviews affect the demand for wine? *American Economic Journal: Applied Economics*, 4(1), 193-211.

Gergaud, O., Storchmann, K., & Verardi, V. (2015). Expert opinion and product quality: Evidence from New York City restaurants. *Economic Inquiry*, 53(2), 812-835.

Ginsburgh, V. (2003). Awards, success and aesthetic quality in the arts. *Journal of Economic Perspectives*, 17(2), 99-111.

Hilger, J., Rafert, G., & Villas-Boas, S. (2011). Expert opinion and the demand for experience goods: an experimental approach in the retail wine market. *Review of Economics and Statistics*, 93(4), 1289-1296.

Moretti, E. (2011). Social learning and peer effects in consumption: Evidence from movie sales. *The Review of Economic Studies*, 78(1), 356-393.

Reinstein, D. A., & Snyder, C. M. (2005). The influence of expert reviews on consumer demand for experience goods: A case study of movie critics. *The Journal of Industrial Economics*, 53(1), 27-51.

Salganik, M. J., Dodds, P. S., & Watts, D. J. (2006). Experimental study of inequality and unpredictability in an artificial cultural market. *Science*, 311(5762), 854-856.

Salganik, M. J., & Watts, D. J. (2008). Leading the herd astray: An experimental study of self-fulfilling prophecies in an artificial cultural market. *Social Psychology Quarterly*, 71(4), 338-355.
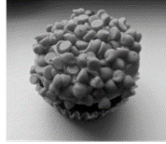
Sorensen, A. T. (2007). Bestseller lists and product variety. *The Journal of Industrial Economics*, 55(4), 715-738.
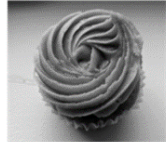
Appendix

Figure A1. Menus in the "Baker's choice" and "Most Popular" treatments

Figure A2. Menus in the "Thumbs-up" recommendation treatments

Figure A3. Location of survey teams around the city circle indicated by stars (left) and aerial picture of one of the festivals showing the city circle with N. 3rd street in the back and Northampton St. on the right (right).

Figure A4. Utility wagon with experiment materials.