Shuffling as a Sales Tactic: An Experimental Study of Selling Expert Advice*

Qichao Shi† James A. Dearden‡ Ernest K. Lai§

February 11, 2021

Abstract

We experimentally investigate the strategic interaction between a product expert and a consumer. The expert privately chooses a ranking methodology to rank two products with uncertain relative merits; the consumer decides whether to acquire the resulting ranking report to guide her product choice. The expert cares only about selling the report; the consumer derives utility from the product itself and an extra ranking attribute controlled by the expert. In equilibrium, the expert chooses sufficiently often a ranking methodology that "shuffles," creating uncertainty in the ranking, to induce the consumer to pay to view the report. The shuffle benefits the expert but could hurt the consumer, which is observed in the laboratory. Consumers are made worse off endogenously when selling reports calls for experts to shuffle more often. With limited field data due to proprietary ranking methodologies, our study provides useful alternative evidence on how profit motives may drive fluctuations of product rankings such as those observed in university rankings.

Keywords: Product Rankings, Expert Advice, Misguidance, Laboratory Experiment

JEL classification: C72; C92; D82; D83; L15

*We are grateful to Weijia (Daisy) Dai, David Goldbaum, and Oliver Yao for valuable comments and suggestions. We also thank seminar and conference participants at Lehigh University, Southwestern University of Finance and Economics, the 2018 China Meeting of the Econometric Society, and the 30th International Conference on Game Theory for helpful comments and discussions. We acknowledge financial support from the Office of the Vice President and Associate Provost for Research and Graduate Studies at Lehigh University (Faculty Research Grant Account No.: 606975).

†School of Economics and China Center for Behavioral Economics and Finance, Southwestern University of Finance and Economics, Chengdu, Sichuan, China. Email: shiqc@swufe.edu.cn.

‡Lehigh University, Department of Economics, Bethlehem, PA 18015, USA. Email: jad8@lehigh.edu.

§Lehigh University, Department of Economics, Bethlehem, PA 18015, USA. Email: kwl409@lehigh.edu.

1 Introduction

Consumers often seek expert advice prior to making purchase decisions. In many cases, the advice takes the form of product rankings. Students and parents consulting university ranking publications (e.g., *U.S. News & World Report Best Colleges Guidebook* and *The Wall Street Journal/Times Higher Education College Rankings*), car buyers acquiring auto rankings (e.g., *Kelley Blue Book New Car Buyer's Guide* and *Consumer Reports New Car Buying Guide*), and home cooks turning to kitchen product rankings (e.g., *Cook's Illustrated*) are but few of the familiar examples. In each of these cases, a ranking publisher collects information on product attributes, chooses a ranking methodology that maps the attributes into a ranking of products, and offers the ranking reports to consumers.

Product rankings guide consumers on what products to choose. The guidance derives from the information provided about product attributes, which otherwise may not be transparent to consumers. Product rankings may also influence consumer choices for reasons unrelated to product information. Owning a highly ranked product may confer sought-after social prestige, not unlike the complementarity between advertisements and advertised products due to imagine concerns (Becker and Murphy, 1993). Consequently, product rankings have the effect of turning a product with n intrinsic attributes into one with $n+1$ attributes, with the extra attribute being the prestige derived from the product's place in the ranking. This in turn suggests that a ranking publisher may influence the values of the ranked products in a way that is extraneous to the intrinsic product attributes.[1] Consumers are willing to pay for a ranking publication not only for informational guidance about existing attributes but also for access to the valuable yet contrived ranking attribute.

Many ranking publishers, including some cited above, are for-profit organizations. The primary objectives of the management of these publishers are nothing short of promoting sales of their ranking reports, boosting subscriptions to their general publications, and maximizing the incidental advertising income. In choosing a ranking methodology, a profit-driven publisher may not necessarily have consumers' best interests in mind, especially when doing so is not aligned with profit motives. An adopted ranking methodology may not rank products in a manner that genuinely reflects the value of the intrinsic product attributes to consumers, and the publisher may leverage the extra attribute its ranking creates to promote sales of its publications. This divergence of interests has been observed by popular press writers. In an article about university rankings that appeared in *The*

[1] Pope (2009) and Luca and Smith (2013) find evidence in, respectively, university and hospital rankings that the rankings themselves, after controlling for product qualities, influence consumer choices. More generally, product reviews, including those submitted by consumers, have been documented to influence consumer decisions ranging from purchases (Chevalier and Mayzlin, 2006; Sun, 2012; Zhu and Zhang, 2010) to returns (Sahoo et al., 2018).

Atlantic, e.g., Tierney (2013) wrote:

> *U.S. News* is always tinkering with the metrics they use, so meaningful comparisons from one year to the next are hard to make. Critics also allege that this is as much a marketing move as an attempt to improve the quality of the rankings: changes in the metrics yield slight changes in the rank orders, which induces people to buy the latest rankings to see what's changed.

The impartiality of a product ranking, if distorted by the publisher's profit motives, would have ramifications not only for consumers but also for other stakeholders of the products being ranked, such as university management in the case of university rankings. Criticisms like the one above are, however, based on indirect evidence that ranking publishers appear to frequently tweak their rankings.[2] There is an inherent difficulty in obtaining direct empirical evidence given that the ranking methodologies of many publishers are proprietary and not transparent to outsiders; by the very nature of the problem, relevant field data are hard to come by. To better understand the incentives and behavior of ranking publishers, who profit from the influences they possess or contrive on consumers, some form of evidence that goes beyond casual observations is needed. In this paper, we provide experimental evidence on how the multiple channels, through which the ranking advice sold by an incentivized expert influences a consumer, interact to shape the expert's adoptions of ranking methods, in a way that may not be in the best interest of the consumer.

We begin by analyzing a ranking-report game, which captures plausible incentives faced by ranking publishers and helps make precise criticisms like the one above. The key insight of our theoretical analysis centers on a phenomenon that we term "shuffling as a sales tactic," which provides an equilibrium rationale for the observation that ranking publishers frequently tweak their rankings and a welfare argument that it is done excessively. In order to generate consumer interests in its publication, a publisher shuffles the ranking outcomes, creating uncertainty over the contrived ranking attribute under the posited prestige effect; a consumer who does not view the ranking report faces the possibility of being deprived of this "ranking value," and this provides an incentive for the consumer to acquire the report.

This behavior is derived from a simple environment suitable for experimentation. A product expert chooses and commits to one of two ranking methods. Each method generates a report ranking two products, whose relative merits based on "intrinsic values," interpreted

[2] On university rankings, the sentiment of criticism is so strong that an entire Wikipedia page is devoted to the issue (`http://en.wikipedia.org/wiki/Criticism_of_college_and_university_rankings_(North_America)`; accessed January 20, 2021). Among the various criticized aspects of university rankings are their year-to-year fluctuations and the seemingly arbitrary ranking formulas. As Tierney (2013) suggests, this tinkering effort by the commercial publishers may have more to do with enhancing sales than with improving the quality of the rankings.

as the values of existing product attributes, are uncertain to a consumer. One ranking method always ranks a particular product first, while the other shuffles and ranks according to the realizations of intrinsic values. Without observing the underlying ranking method, the consumer decides whether to acquire and view the costly ranking report and then chooses a product. The expert cares only about selling the report. The consumer obtains the intrinsic value of the chosen product and an additional ranking value if the product is ranked first. "Shuffling as a sales tactic" manifests as a mixed-strategy equilibrium in which the expert chooses the rank-shuffling method sufficiently often to induce the consumer to acquire the report. Remarkably, this happens even when the ranking method does not always provide useful informational guidance and the consumer is worse off under it, providing a sense that the shuffle of the rankings is excessive.

This last scenario drives our treatment design and forms the focus of our experimental inquiry. In our two main treatments, the treatment parameters are selected so that the rank-shuffling method does not always guide the consumer to choose the most intrinsically valuable product—it sometimes misguides.[3] In the expert-preferred equilibria in which ranking reports are acquired, the consumer sometimes ends up with a product that is inferior in terms of intrinsic value, and that is made up by the "artificial" ranking value. The two treatments differ by an intrinsic-value parameter. The treatment variation results in different degrees of "misguidance" in the event that the consumer follows a shuffled report to purchase an intrinsically inferior product. The higher the degree of misguidance in this event, the lower the marginal gain from acquiring the report. Equilibrium then requires the expert to choose the rank-shuffling method more often to maintain the expected marginal gain sufficient to lure the consumer to pay. The result is that the expert shuffles more often when doing so hurts the consumer more.

Our experimental findings support the comparative statics. The experts in the treatment with potentially more misguiding rank-shuffling method chooses the method with higher aggregate frequency, which is corroborated by an individual panel data analysis controlling for consumers' report acquisitions. Responding to the sentiment that motivates our study, we indeed find that by shuffling the rankings experts benefit at the expense of consumers; comparing payoffs across treatments, we find that, when the rank-shuffling method is associated with a higher degree of misguidance, consumers are *endogenously* made worse off due to more frequent adoptions of the method, while experts are better off.

Related Literature. Our study is related to two separate strands of literature. In terms of the subject matter, we contribute to the literature on product rankings and more generally non-seller-provided product information. In terms of the experiment, our study is broadly

[3] There is an immense equilibrium multiplicity in our game. We also include a third, exploratory treatment with parameters under which the equilibria allow for a wider range of possibilities.

related to the literature on experimental signaling games.

In respect of product rankings, our game is adapted from the more general model of university rankings analyzed in Dearden et al. (2019). In their multi-period model, a finite number of universities with a finite number of attributes are ranked in each period by a ranking publisher, whose per-period payoff depends on the number of students who view the ranking. The publisher's strategy entails choosing the weight of each attribute for an "attribute-and-aggregate" ranking methodology. We develop a static version of the model, simplified for experimentation yet capturing their key result: the prestige effect of university rankings incentivizes the publisher to add more uncertainty into the ranking than is optimal for the students. Their dynamic model also yields insight that is absent in our static setting. They find that as the prestige effect grows over time, the publisher gradually moves away from the ranking method that best serves students by adding noise.

The prevalence of random product recommendations in the presence of prestige effects is also observed in the fashion industry. Editors of influential fashion magazine determine a season's "it" products in a seemingly random manner, in effect offering non-seller-based product recommendations that fluctuate season by season. Consumers who follow the recommendations derive social prestige from signaling to others that they are "in the club."[4] Kuksov and Wang (2013) analyze a model of fashion that shares a common theme with ours. Stylish consumers, who prefer to be identified, have exclusive access to a coordinator interpreted as a fashion magazine that makes product recommendations. By following the magazine's recommendations, which are random in nature, these high-type stylish consumers separate themselves from the low types, thus maximizing their utility under the prestige effect. The seeming randomness in product rankings in our case and fashion hits in theirs are commonly rationalized as outcomes of maximizing behavior.

Product rankings or recommendations by product magazines are not the only non-seller sources of product information—online product reviews are another. While these reviews submitted by consumers, e.g., physician ratings (Lu and Rui, 2018), have been shown to provide useful product information, the ubiquity of fake reviews, either submitted by sellers themselves or competitors (Mayzlin et al., 2014), dilutes the value of information provided by customer review platforms (Anderson and Simester, 2014), to the extent that some platforms have resorted to algorithms to filter out suspicious reviews (Luca and Zervas, 2016). Our research contributes to the picture of non-seller-based product information by illustrating that ranking publishers who take no interests in consumer choices may still have incentives to provide misguiding product information under their own profit motives.

[4]Fashion items, and to a certain extent universities, are examples of status goods. The value of a status good depends not only on product attributes but also on who are the consumers—the caliber of the "club members" (Kuksov and Xie, 2012).

On the experimental front, our experiment shares a common feature with those on signaling games. Broadly defined, games with signaling opportunities encompass environments in which private information is transmitted via payoff-dependent (costly) messages (Spence, 1973), payoff-independent (cheap-talk) messages (Crawford and Sobel, 1982), and state-dependent (verifiable) messages (Grossman, 1981; Milgrom, 1981). Some of the early experimental work includes, e.g., Brandts and Holt (1992) and Banks et al. (1994) for costly signaling, Dickhaut et al. (1995) and Gneezy (2005) for cheap talk, and Forsythe et al. (1989) and King and Wallin (1991) for verifiable disclosures. Similar to these experiments, ours features a sender of information who influences the action of a receiver via the message sent; the ranking report in our game can be viewed as a message.

There are nevertheless a number of critical differences. While it is costless for the expert in our game to generate the ranking report, the report directly influences the consumer's payoff by means of the ranking value. This feature makes our environment not readily fit into the three canonical message categories. A more important distinction is that in our game information is not transmitted as a direct execution of the expert's strategy; rather, it is transmitted under a signaling rule (ranking method), a mapping from states to messages, that the expert commits to use. The "private information" in our game is not the exogenous state of the world but the endogenous mapping chosen by the expert, and there is no channel in place for the expert to reveal this information.[5]

The fact that our expert's choice of a ranking method amounts to committing to a signaling rule also relates our study to the more recent literature on Bayesian persuasion (Kamenica and Gentzkow, 2011), with experimental attempts including Nguyen (2016), Au and Li (2018), and Fréchette et al. (2020). Despite the similarity in the generation of information, there is, by contrast, no element of selling information in Bayesian persuasion.

The remainder of the paper proceeds as follows. Section 2 presents and analyzes our experimental ranking-report game. Section 3 describes our experimental design and hypotheses. We report our laboratory findings in Section 4. Section 5 concludes.

[5] It is worth comparing the roles played by uncertainty or randomness in our game and in cheap-talk games. Theory (Krishna and Morgan, 2004; Blume et al., 2007; Goltsman et al., 2009) and experiment (Blume et al., 2020) have shown that random transmissions of messages could improve cheap-talk communication, resulting in Pareto improvements. In our game, randomness has a different welfare effect, in which it benefits one party but could hurt the other. Another contrast is that those welfare-improving randomly transmitted messages in cheap-talk games are an exogenous property of the communication process, while in our case the random ranking reports are consequences of endogenous choices.

2 The Ranking-Report Game

2.1 The Setup

There are two players, a product expert (he) and a consumer (she), and two products, A and B. The expert chooses a ranking *method* to rank the two products and sells the resulting ranking *report* to the consumer. The imperfectly informed consumer makes two decisions, whether to acquire the ranking report and which product to purchase.

Consumer Utility. The consumer derives utility from the *intrinsic attributes* and the *ranking attribute* of the purchased product. The intrinsic attributes, which may include quality, features, price, etc., are determined exogenously by the manufacturer not being modeled. The ranking attribute, which concerns how the product is ranked, is determined endogenously by the expert.

A product's intrinsic attributes give rise to its *intrinsic value*, v. The intrinsic value of Product A, $\bar{v}_A > 0$, is fixed and commonly known. The intrinsic value of Product B, v_B, is uncertain, taking one of two values, 0 and $\bar{v}_B > 0$, not observed by the consumer. The common prior is that $v_B = \bar{v}_B$ with probability $0 < p < 1$. We assume that $\bar{v}_A > p\bar{v}_B$ so that Product A is intrinsically more valuable *ex ante*.[6] The ranking attribute of a product yields to the consumer either $\gamma > 0$ if the product is ranked first or 0 otherwise, independent of its intrinsic value. We call γ the *ranking value* of the first-ranked product.

Product rankings serve comparable functions as advertising. The way we model the influence of the expert's ranking on consumer utility is consistent with two economic views on why consumers respond to advertising (see, e.g., Bagwell, 2007). The first view considers advertising as information, through which consumers learn about product attributes. In our game, the ranking report may provide information about v_B. The second view considers advertising as a good that is complementary to the consumption of the advertised product, where an increase in advertising raises the marginal utility of the product. Analogously, in our game an increase in the ranking of a product raises the utility of the product by γ.[7]

[6] For a concise exposition, we also assume that $\bar{v}_A \neq \bar{v}_B$ so that we do not need to refer to the uninteresting case where the two products value intrinsically the same *ex post*.

[7] The complementary view of advertising (Becker and Murphy, 1993) suggests that consumers have imagine concerns and consuming an advertised product gives rise to valuable "social prestige." In the case of product rankings, a similar prestige may be associated with consuming a highly ranked product. The complementarity may also result from the future economic value conferred by the rankings. For example, a car model that is ranked top in a popular auto ranking is likely to have a higher resale value. Note also that the literature on advertising holds a third view of advertising, considering it as persuasive aimed at directly altering consumer taste. In our game, the intrinsic and the ranking values are additively separable, and the ranking report does not directly alter the intrinsic value; the influence of the ranking report on consumer utility is more in line with the complementary view than with the persuasive view.

Ranking Methods. The expert chooses and commits to a ranking method, learns the realized intrinsic value of Product B, and then issues a ranking report according to the adopted method. A ranking method is a mapping, $r : \{0, \bar{v}_B\} \to \{A, B\}$, which specifies for each possible intrinsic value of Product B a ranking report available for the consumer. Report A (B) indicates that Product A (B) is ranked first.

There are two choices of ranking method, r_A and $r_{A\|B}$. Method r_A always ranks Product A first, i.e., $r_A(0) = r_A(\bar{v}_B) = A$. Method $r_{A\|B}$ ranks the products according to the intrinsic value of Product B, as follows:

$$r_{A\|B}(v_B) = \begin{cases} A & \text{if } v_B = 0, \\ B & \text{if } v_B = \bar{v}_B. \end{cases}$$

The consumer can access the ranking report if and only if she pays a fee $c > 0$. To capture the proprietary nature of ranking methodologies in practice, we assume that the consumer does not observe the adopted ranking method, whether acquiring the report or not. With or without viewing the report, the consumer then chooses between Products A and B.[8]

Our experiment explores the interplay of two properties of ranking methods, one about *product guidance* and the other *uncertainty* over the ranking value. A ranking report is said to provide product guidance if, given the realization of v_B, it ranks the more intrinsically valuable product first. A ranking method is *misguidance-proof* if it never generates a misguiding report. For $r_{A\|B}$, its report A always provides guidance since $\bar{v}_A > 0$, while its report B provides guidance only if $\bar{v}_A < \bar{v}_B$; if $\bar{v}_A > \bar{v}_B$, report B is misguiding, and thus $r_{A\|B}$ is non-misguidance-proof. The property is mutually exclusive; for any given $\bar{v}_A$ and $\bar{v}_B$, one and only one of the ranking methods is misguidance-proof.[9]

While a ranking report provides information about the ranking value γ, a ranking method may create uncertainty about it. This second property of a ranking method contributes to determine the demand for its report. Knowledge of the adoption of the *rank-shuffling* $r_{A\|B}$ is not sufficient for the consumer to learn which product is ranked first, thus generating a demand for its report even when the consumer correctly anticipates the expert's choice of method. On the other hand, knowledge of the adoption of the *rank-invariant* r_A renders viewing its report to learn which product carries γ unnecessary.

The concepts of misguidance-proofness and rank shuffling constitute the two building blocks of our formalization of the idea that ranking publishers excessively inject uncertainty

[8]Product prices are assumed to be constituents of the intrinsic values so that the consumer does not explicitly pay for the products.

[9]When $r_{A\|B}$ is non-misguidance-proof, the inequality $\bar{v}_A > \bar{v}_B$ implies that r_A is misguidance-proof because Product A is always intrinsically more valuable. On the other hand, when $r_{A\|B}$ is misguidance-proof (i.e., $\bar{v}_A < \bar{v}_B$), the report A generated by r_A is misguiding with probability p—it ranks Product A first even when Product B has a higher intrinsic value—and this renders r_A non-misguidance-proof.

into product rankings. As will be elaborated below, the case of particular interest, which forms the focus of our experiment and accentuates the tension between the expert's profit motive and the consumer's interests, is when the rank-shuffling $r_{A\|B}$ is non-misguidance-proof; by choosing $r_{A\|B}$ rather than r_A when $\bar{v}_A > \bar{v}_B$, the expert creates uncertainty with respect to γ and, in the event that $v_B = \bar{v}_B$, misguides rather than guides the consumer.

Timing and Payoffs. We summarize the sequences of moves through specifying players' strategies. The expert chooses between r_A and $r_{A\|B}$ before observing the realized v_B. A mixed strategy of the expert, $s \in [0,1]$, specifies the probability that he chooses $r_{A\|B}$. A behavior strategy of the consumer consists of two components: (a) a report-acquisition decision, $q \in [0,1]$, specifying the probability that she acquires the ranking report, and (b) a product-purchasing rule, $a : \{\varnothing, A, B\} \to [0,1]$, specifying for each possible report, A, B, or none ($\varnothing$), a probability that she purchases Product A. We call $a(\varnothing)$ the consumer's *default product choice*. Note that the consumer's report-acquisition decision and the expert's choice of ranking method are strategically simultaneous.

The expert's payoff consists of his revenue from selling the ranking report. His choice of ranking method does not affect his payoff, nor does he take interest in the product choice of the consumer. He earns $\pi > 0$ if the consumer acquires the report.[10]

The consumer's payoff equals her utility from consuming the product of her choice, which consists of the intrinsic value and any ranking value, minus the ranking report fee should she acquire one; a consumer who chooses product $K \in \{A, B\}$ with intrinsic value v_K, has report-acquisition decision outcome $q_o \in \{\text{acquire}, \text{not acquire}\}$, and faces an expert issuing ranking report $r(v_B)$ receives a payoff of

$$u(K, v_K, q_o, r(v_B)) = v_K + \gamma \mathbb{I}_{r(v_B)}(K) - c\mathbb{I}_{ACQ}(q_o),$$

where $\mathbb{I}_{r(v_B)}(K)$ is an indicator function taking the value of 1 if $K = r(v_B)$, i.e., Product K is ranked first, and 0 otherwise, and $\mathbb{I}_{ACQ}(q_o)$ is an indicator function taking the value of 1 if q_o = acquire and 0 otherwise. Note that the consumer receives γ for choosing the first-ranked product irrespective of her report-acquisition decision.

[10] We allow the expert's revenue to be different from the report fee the consumer pays, specifically considering that $\pi > c$. In practice, ranking publishers typically have additional income sources incidental to the sales of ranking reports, such as advertising. Alternatively, our two-player game can be considered as a reduced-form model of one expert selling ranking reports to multiple identical consumers, in which case π would be a multiple of c.

2.2 Equilibrium Characterization

We begin by highlighting some useful facts about how product guidance and ranking value interact to determine the influences of the ranking methods on the consumer. A ranking method is said to be *influential* if all the reports that may be generated by the method influence the consumer to choose the first-ranked products. Since r_A only generates report A, the method is influential if report A influences the consumer to choose Product A. For $r_{A\|B}$ to be influential, both reports A and B are required to influence the consumer.

While we assume that $\gamma > 0$, to bring out its role in determining the influence of a ranking method we speak of the absence of ranking value ($\gamma = 0$) in the following two facts:

Fact 1. *Method r_A is always influential, even in the absence of ranking value.*

While r_A only generates report A, report A can also be generated by $r_{A\|B}$, which occurs when the intrinsic value of Product B is 0. Upon seeing that Product A is ranked first, the consumer assigns, relative to the prior, a lower probability that $v_B = \bar{v}_B$. Fact 1 follows from the assumption that Product A is intrinsically more valuable *ex ante* ($\bar{v}_A > p\bar{v}_B$).

Fact 2. *When $r_{A\|B}$ is misguidance-proof, it is always influential, even in the absence of ranking value. When $r_{A\|B}$ is not misguidance-proof, it is influential only if $\gamma > 0$.*

The case for the influence of report A generated by $r_{A\|B}$ is the same as that for Fact 1. For report B, since it is exclusive to $r_{A\|B}$, when the consumer sees that Product B is ranked first, she is certain that its intrinsic value is $\bar{v}_B$. There are two cases. If $\bar{v}_B > \bar{v}_A$ so that report B provides guidance, the consumer follows to choose Product B even without regard to the ranking value. If $\bar{v}_A > \bar{v}_B$ so that report B is misguiding, she will not follow the report if there is no ranking value accompanying the purchase of the first-ranked product.

The consumer will not pay for the ranking report if it is generated by a non-influential ranking method. Fact 2 highlights the fact that the information provided by $r_{A\|B}$ about the intrinsic value of Product B *by itself* renders the ranking method influential, generating possible demand for its report, only if it always provides product guidance. If it does not, the expert will have to rely on the ranking value to generate demand. We fully develop this observation by analyzing the perfect Bayesian equilibria of the game, in which players best respond to beliefs and beliefs are formed via Bayes' rule whenever possible. We denote by $\mu_{r_{A\|B}}$ the consumer's *strategy belief* that the expert chooses $r_{A\|B}$ and $\mu_{\bar{v}_B}$ her *product belief* that the intrinsic value of Product B is $\bar{v}_B$.[11]

[11] Although the expert's choice of ranking method and the consumer's report-acquisition decision are strategically simultaneous, a perfect Bayesian equilibrium explicitly specifies the consumer's strategy beliefs, in addition to her product beliefs.

We first characterize the consumer's optimal product-purchasing rule, starting with the case where she has acquired the ranking report (all proofs are relegated to Appendix A):

Lemma 1. *The optimal product choices of a consumer who has acquired the ranking report are as follows:*

(a) if $\gamma > \bar{v}_A - \bar{v}_B$, the consumer's unique best response is to purchase the first-ranked product ($a(A) = 1$ and $a(B) = 0$);

(b) if $\gamma < \bar{v}_A - \bar{v}_B$, the consumer's unique best response is to purchase Product A irrespective of the report ($a(A) = a(B) = 1$); and

(c) if $\gamma = \bar{v}_A - \bar{v}_B$, the consumer's unique set of best responses consists of purchasing the first-ranked Product A ($a(A) = 1$) and randomizing between the first-ranked Product B and second-ranked Product A with any probability ($a(B) \in [0, 1]$).

Lemma 1 concerns the parts of the consumer's product-purchasing rule $a(A)$ and $a(B)$. In all cases, $a(A) = 1$, a restatement of Fact 1. For the product choice after the consumer views report B, $a(B) = 0$ in case (a), $a(B) = 1$ in case (b), and $a(B) \in [0, 1]$ in case (c). In addition to the always satisfied $a(A) = 1$, $r_{A\|B}$ is influential if and only if we also have that $a(B) = 0$. The lemma adds to Fact 2 by specifying the parameter profiles of γ, $\bar{v}_A$, and $\bar{v}_B$ under which $r_{A\|B}$ is or is not influential.

We next characterize the consumer's optimal default product choice. Unlike the case where the consumer has acquired the report, the expert's strategy s plays a role here:

Lemma 2. *The optimal default product choice of a consumer who has not acquired the ranking report is as follows:*

(a) if $\bar{v}_A - p\bar{v}_B > (2p - 1)\gamma$, the consumer's unique best response is to purchase Product A ($a(\varnothing) = 1$); and

(b) if $\bar{v}_A - p\bar{v}_B \leq (2p - 1)\gamma$, there exists $\tilde{s} = \frac{\bar{v}_A - p\bar{v}_B + \gamma}{2p\gamma} \in (0, 1]$ such that (i) for $s \in [0, \tilde{s})$, the consumer's unique best response is to purchase Product A ($a(\varnothing) = 1$), (ii) for $s = \tilde{s}$, randomizing between the two products with any probability ($a(\varnothing) \in [0, 1]$) constitutes a best response, and (iii) for $\tilde{s} < 1$ and $s \in (\tilde{s}, 1]$, her unique best response is to purchase Product B ($a(\varnothing) = 0$).

Without the report, the consumer's product belief $\mu_{\bar{v}_B}$ equals the prior p, and correct beliefs required by equilibrium dictate that her strategy belief $\mu_{r_{A\|B}}$ coincides with the expert's strategy s. In case (a), $a(\varnothing) = 1$ irrespective of s. To see the significance of the condition $\bar{v}_A - p\bar{v}_B > (2p-1)\gamma$, note that even though Product A is intrinsically more valuable *ex ante*, unless the expert always chooses r_A ($s = 0$), Product B may be ranked first and represent a better choice because of the ranking value γ; the condition, which states that the excess in expected intrinsic value of Product A over Product B outweighs the excess in expected ranking value of Product B over Product A in the most unfavorable case for choosing Product A ($s = 1$), guarantees that Product A is always preferred. When the condition does not hold, the consumer may still, as in case (b), prefer Product A if the expert chooses r_A sufficiently often, otherwise she would prefer Product B or be indifferent.

We proceed to characterize the equilibria under parameter profile $\gamma > \bar{v}_A - \bar{v}_B$, which corresponds to the meaningful case where both ranking methods are influential and the consumer acquires the ranking report under some circumstances.[12] The following proposition describes what these circumstances, defined in terms of the expert's strategies, are:

Proposition 1. *The perfect Bayesian equilibria of the ranking-report game in which the consumer adopts the product-purchasing rule* $a(\varnothing) = a(A) = 1$ *and* $a(B) = 0$ *are as follows: for threshold* $\bar{s} = \frac{c}{p(\bar{v}_B - \bar{v}_A + \gamma)} \in (0,1)$*, there are three sets of equilibria, in which*

(a) the expert chooses $r_{A\|B}$ *with probability* $s \in [0, \bar{s})$*, and the consumer does not acquire the ranking report* ($q = 0$)*;*

(b) the expert chooses $r_{A\|B}$ *with probability* $s \in (\bar{s}, 1]$*, and the consumer acquires the ranking report* ($q = 1$)*; and*

(c) the expert chooses $r_{A\|B}$ *with probability* $\bar{s}$*, and the consumer randomizes with any probability* ($q \in [0,1]$)*.*

For the equilibria in which $a(A) = 1$*,* $a(B) = 0$*, and* $a(\varnothing) \in [0,1)$*, there are two cases: (i) if* $\bar{v}_A + \gamma < \bar{v}_B$*, (a)–(c) with* $\bar{s} \in (0,1)$ *replaced by* $\hat{s} = \frac{\bar{v}_A - p\bar{v}_B + \gamma - c}{p(\bar{v}_A + \gamma - \bar{v}_B)} \in (0,1)$ *constitute the equilibria; and (ii) if* $\bar{v}_A + \gamma \geq \bar{v}_B$*, (a)–(c) with* $\bar{s} \in (0,1)$ *replaced by* $\hat{s} \in (0,1)$ *and the consumer instead choosing* $q = 1$ *in (a) and* $q = 0$ *in (b) constitute the equilibria. For* $\bar{s} \geq 1$ *or* $\hat{s} \notin (0,1)$*, maintaining* $s \in [0,1]$ *rules out, for each case, one or two sets of strategy profiles as equilibria.*

There are uncountably many equilibria. Their structure can be elucidated by examining the consumer's expected payoffs at the strategically simultaneous stage, where the expert chooses the ranking method and the consumer makes the report-acquisition decision. Table 1 presents the case where Product A is the default product choice.

[12] The condition in Lemma 1(b), $\gamma < \bar{v}_A - \bar{v}_B$, is associated with the uninteresting case where the consumer never acquires any ranking report. The consumer does acquire the ranking report under some circumstances when the condition in Lemma 1(c), $\gamma = \bar{v}_A - \bar{v}_B$, holds, but for brevity we also exclude this case. The characterizations and proofs for these two cases are available from the authors upon request.

Table 1: Consumer's Expected Payoffs: Product A as Default Product

		Report-Acquisition Decision		
		(c) $q = 0$	(d) $q = 1$	(d)−(c)
Ranking Method	(a) r_A	$\bar{v}_A + \gamma$	$\bar{v}_A + \gamma - c$	$-c$
	(b) $r_{A\|\|B}$	$\bar{v}_A + (1-p)\gamma$	$p\bar{v}_B + (1-p)\bar{v}_A + \gamma - c$	$p(\bar{v}_B - \bar{v}_A + \gamma) - c$
	(b)−(a)	$-p\gamma$	$-p(\bar{v}_A - \bar{v}_B)$	

When the consumer purchases Product A with or without viewing report A, the (only) report that is generated by r_A is superfluous; acquiring the report under r_A brings no marginal gain but a fee, resulting in a *net* marginal loss for the consumer, $-c < 0$.[13] In equilibrium, therefore, the consumer would not acquire the report when the expert chooses r_A. For $r_{A\|B}$, in addition to report A it also generates, with probability p, report B. Unlike report A, viewing report B makes a difference—the consumer receives $\bar{v}_B + \gamma$ if viewing it and $\bar{v}_A$ otherwise. The consumer thus derives an expected marginal gain from viewing the report generated by $r_{A\|B}$. If the report fee c is less than this expected gain, acquiring the report under $r_{A\|B}$ brings a net marginal gain $p(\bar{v}_B - \bar{v}_A + \gamma) - c > 0$. In equilibrium, therefore, the consumer acquires the report when the expert chooses $r_{A\|B}$.

The first threshold in Proposition 1, $\bar{s} = \frac{c}{p(\bar{v}_B - \bar{v}_A + \gamma)}$, is the point where the probability of $r_{A\|B}$ being chosen renders the consumer indifferent; should the expert's strategy deviates from this threshold, the consumer strictly prefers to either acquire or not to acquire the report.[14] As will be further discussed below, our main experimental treatments explore a comparative statistics in relation to $\bar{s}$: the threshold increases as the difference $\bar{v}_B - \bar{v}_A$ decreases. The net marginal gain from acquiring the report under $r_{A\|B}$ decreases as $\bar{v}_B - \bar{v}_A$ decreases, while acquiring the report under r_A results in a constant loss. To induce the consumer to acquire the report for a smaller $\bar{v}_B - \bar{v}_A$, it takes a more frequent choice of $r_{A\|B}$ to make up for the smaller gain.

[13]The column "(d)−(c)" in Table 1 (and the forthcoming Table 2) contains, for given ranking methods, the consumer's payoff differences between acquiring and not acquiring the report, which take into account the report fee and are construed as the *net* marginal gain/loss from *acquiring* the report. The payoff differences excluding the report fee are referred to as the marginal gain from *viewing* the report.

[14]Our parameter restriction for the first-ranked products to be chosen, $\gamma > \bar{v}_A - \bar{v}_B$, implies that $p(\bar{v}_B - \bar{v}_A + \gamma) > 0$, which in turn implies that $\bar{s} > 0$. The case where $\bar{s} \geq 1$ corresponds to the situation where the report fee is higher than or equal to the marginal gain from acquiring the report under $r_{A\|B}$, and in equilibrium the consumer does not acquire the report irrespective of the expert's strategy.

Table 2: Consumer's Expected Payoffs: Product B as Default Product

		Report-Acquisition Decision		
		(c) $q = 0$	(d) $q = 1$	(d)−(c)
Ranking Method	(a) r_A	$p\bar{v}_B$	$\bar{v}_A + \gamma - c$	$\bar{v}_A - p\bar{v}_B + \gamma - c$
	(b) $r_{A\|\|B}$	$p(\bar{v}_B + \gamma)$	$p\bar{v}_B + (1-p)\bar{v}_A + \gamma - c$	$(1-p)(\bar{v}_A + \gamma) - c$
	(b)−(a)	$p\gamma$	$-p(\bar{v}_A - \bar{v}_B)$	

Table 2 presents the consumer's payoffs under the alternative default product choice of Product B. The characteristics of the two reports are now reversed, with report B becoming superfluous while report A makes a difference. Since report A is generated under both r_A and $r_{A||B}$, viewing the report under both ranking methods each brings a marginal gain. If $\bar{v}_A + \gamma < \bar{v}_B$, the marginal gain from viewing the report under $r_{A||B}$ is larger than that under r_A; for levels of c that are between these two marginal gains, acquiring the report under r_A results in a net marginal loss $\bar{v}_A - p\bar{v}_B + \gamma - c < 0$, while that under $r_{A||B}$ gives a net marginal gain $(1-p)(\bar{v}_A + \gamma) - c > 0$. In equilibrium, therefore, the consumer acquires the report under $r_{A||B}$ but not under r_A. The opposite case prevails if $\bar{v}_A + \gamma \geq \bar{v}_B$. The second threshold in Proposition 1, $\hat{s} = \frac{\bar{v}_A - p\bar{v}_B + \gamma - c}{p(\bar{v}_A + \gamma - \bar{v}_B)}$, is the point where the consumer is indifferent between acquiring and not acquiring the report when Product B is the default.[15]

For the expert's equilibrium strategy, since he cares only about selling the report and is otherwise indifferent in his choice of ranking method and the consumer's product choice, he is willing to randomize between the two ranking methods. The randomization is restricted only to support a particular report-acquisition decision of the consumer in equilibrium.

2.3 Shuffling as a Sales Tactic: A Discussion

Among the many equilibria, the expert prefers those in which he sells the ranking report. While this propensity of the expert is simple and clear, how he induces a willingness to pay for his report brings to the fore some important interpretations of equilibria. The full

[15]Unlike $\bar{s}$, which is always positive, we can have that $\hat{s} \leq 0$, which corresponds to the situations where the report fee is, for $\bar{v}_A + \gamma < \bar{v}_B$ ($\bar{v}_A + \gamma \geq \bar{v}_B$), less than (higher than) or equal to any convex combination of the two marginal gains so that in equilibrium the consumer acquires (does not acquire) the report irrespective of the expert's strategy. The situations are reversed ("less than/acquires" and "higher than/does not acquire" are interchanged) for the case where $\hat{s} \geq 1$.

characterization contains equilibria in which the expert sells under either ranking method. Those in which the rank-shuffling $r_{A\|B}$ is chosen sufficiently often for sales to take place are of particular interests. They form the focus of our experimental inquiry, guide the design of our major treatments, and warrant further discussion.

A key observation on how the expert generates interests in his report is that the uncertainty over the ranking value, which the consumer pays to avoid, can be viewed as a creation of the expert. By choosing the rank-shuffling $r_{A\|B}$ rather than the rank-invariant r_A, it is as if the expert peddled a solution for a problem he created. The pure-strategy equilibrium in which the consumer acquires the report under a *non-misguidance-proof* $r_{A\|B}$, i.e., when $\bar{v}_A > \bar{v}_B$, provides a lucid illustration. It shows how shuffling the ranking could serve as a sales tactic for the expert to the detriment of the consumer.

To be sure, the fact that it is an equilibrium means that acquiring the report is the consumer's best course of action given the expert's choice of $r_{A\|B}$. The consumer's willingness to pay for the report derives from the expected marginal gain from viewing it. For $r_{A\|B}$ with Product A being the default, it is $p(\bar{v}_B - \bar{v}_A + \gamma)$. This willingness to pay can be decomposed into two parts: one for product guidance, $p(\bar{v}_B - \bar{v}_A)$, and the other for the resolution of uncertainty over the ranking value, $p\gamma$. Since $\bar{v}_A > \bar{v}_B$, the willingness to pay for product guidance is negative, i.e., the consumer needs to be compensated for the misguidance to acquire the report. In equilibrium, her willingness to pay for the resolution of uncertainty is enough to cover that; she locks in the ranking value at the expense of paying the report fee and sometimes settling with a less intrinsically valuable product.

In terms of welfare, however, the consumer is worse off under $r_{A\|B}$ than under r_A, and that is true not only when she acquires the report but also when she does not; from the perspective of the consumer's welfare, the shuffling by the expert is excessive. The row "(b)−(a)" in Table 1, which presents the consumer's payoff differences under the two ranking methods, makes this clear. Relative to her *no-report* payoff under r_A, not viewing the report under $r_{A\|B}$ deprives the consumer of the ranking value γ with probability p. Relative to her *with-report* payoff under r_A, viewing the report under $r_{A\|B}$ also strips the consumer of the difference between $\bar{v}_A$ and $\bar{v}_B$ with probability p. Given that $\gamma > \bar{v}_A - \bar{v}_B$, she is nevertheless more worse off without the report. By shuffling the ranking even when there is no guidance value in doing so, it is as if the expert created two unfavorable situations for the consumer and profited from offering the less unfavorable situation with a price tag.[16]

[16] The discussion accentuates the tension between the expert's payoff motive and the consumer's welfare by focusing on the case where the consumer acquires the report under a non-misguidance proof $r_{A\|B}$. According to Proposition 1, Product A must be the default product in this case. The rank shuffling serves the same purpose of enhancing the consumer's willingness to pay when $r_{A\|B}$ is misguidance-proof. The main difference is that the consumer locks in the ranking value and *always* obtains product guidance under $r_{A\|B}$ when it is misguidance-proof. Another difference is that Product B may be the default in this case.

3 Experimental Design

3.1 Treatments and Procedures

Guided by the equilibrium characterizations, we design treatments by assigning values to the six parameters in the game, $\bar{v}_A$, $\bar{v}_B$, γ, p, c, and π. Our primary focus is on the equilibria in which the expert shuffles to sell at the consumer's expense; our two main treatments feature a non-misguidance-proof $r_{A||B}$, with Product A being the default choice. The treatment variation is the extent of misguidance. We also include an additional treatment that explores the richer comparative statics when either product could be the default.

All three treatments share the same parameter values of $\bar{v}_A = 100$, $\gamma = 250$, $p = 0.6$, $c = 70$, and $\pi = 350$, and we vary the value of $\bar{v}_B$ at 60 and 80 for the main treatments and 120 for the exploratory treatment.[17] Except for the prior probability p, the parameter values are induced as monetary amounts that subjects receive or pay.[18] The three sets of parameters commonly satisfy the restrictions that (a) Product A is intrinsically more valuable *ex ante* ($\bar{v}_A > p\bar{v}_B$), and (b) purchasing the first-ranked product after viewing the report is the unique best response ($\gamma > \bar{v}_A - \bar{v}_B$).

Table 3: Experimental Treatments

	NMP-SHUF-H	*NMP-SHUF-L*	*NMP-INVA*
$\bar{v}_B$	60	80	120

Note: All three treatments share the following parameter values: $\bar{v}_A = 100$, $\gamma = 250$, $p = 0.6$, $c = 70$, and $\pi = 350$.

The two main treatments with $\bar{v}_B = 60$ and $\bar{v}_B = 80$, which we name *NMP-SHUF-H* and *NMP-SHUF-L* respectively, further satisfy that (a) irrespective of the expert's strategy, Product A is the default product $[\bar{v}_A - p\bar{v}_B > (2p-1)\gamma]$, and (b) $r_{A||B}$ is non-misguidance-proof ($\bar{v}_A > \bar{v}_B$). Note that $\bar{v}_A - \bar{v}_B$, which equals 40 for $\bar{v}_B = 60$ and 20 for $\bar{v}_B = 80$, measures the extent of misguidance, which is reflected in the treatment names. For example, *NMP-SHUF-H* captures the fact that the featured non-misguidance-proof (*NMP*) rank-shuffling $r_{A||B}$ (*SHUF*) has a high (*H*) extent of misguidance. The additional treatment with $\bar{v}_B = 120$ satisfies that (a) depending on the expert's strategy, either product can be the default choice $[\bar{v}_A - p\bar{v}_B \leq (2p-1)\gamma]$, and (b) $r_{A||B}$ is misguidance-proof ($\bar{v}_A < \bar{v}_B$). We name this treatment *NMP-INVA*; in this case the rank-invariant r_A (*INVA*) is non-misguidance-proof. Table 3 summarizes the parameter values adopted for the treatments.

[17] In reference to footnote 10, we thus have $\pi = 5c$ in our design of the treatment parameters.

[18] Other than the equilibrium characterization, another criterion in our choices of parameters is salience: the values are chosen so that the payoff differences from different choices are reasonably large.

The experiment was conducted using *oTree* (Chen et al., 2016) at the Financial Services Lab of Lehigh University. A total of 148 subjects, who had no prior experience with the experiment, participated. Upon arrival at the laboratory, subjects were instructed to sit at individual computer terminals separated by partitions. Each subject received a copy of the experimental instructions. The instructions were read aloud, aided by slide illustrations. A comprehension quiz and a practice round followed.[19]

Nine experimental sessions, three for each treatment, were conducted using a *between-subject* design. Each session was participated by 14 to 20 subjects, with half of them randomly assigned to the role of an expert and the other half to the role of a consumer. Roles remain fixed throughout a session. Experts and consumers in a session were *randomly matched* in each round to form groups of two to play *40 rounds* of the game. They were presented with two Products, A and B. The value of Product A was fixed at 100. The value of Product B was random, taking the value of either 0 or, using *NMP-SHUF-L* as an example, 80. Subjects were told that the computer selected 80 with 60% chance.

In each round, the expert made one decision, choosing in a neutral frame between Ranking Method 1 (r_A) and Ranking Method 2 ($r_{A||B}$). The consumer made two decisions, first deciding whether to acquire the ranking report and then choosing a product. In the case where the consumer decided to acquire the report, she would receive, before choosing a product, a report ("Product K is ranked first") generated according to the randomly selected value of Product B and the chosen ranking method. In the case where the consumer decided not to acquire the report, she would proceed directly to choose a product. A report not observed by the consumer would still be generated to determine the ranking value.

Each round ended with an information feedback, which summarized the events in the round, including the ranking method chosen by the expert, the randomly selected value of Product B, the generated ranking report, the consumer's report-acquisition decision, the consumer's product choice, and the reward.

We randomly selected three out of the 40 rounds for calculating subject payments. The average reward a subject earned in the three selected rounds was converted into US Dollars at a fixed and known exchange rate of US$1 per 20 reward points. A show-up fee of US$5 was also paid. Subjects on average earned US$14.75. A session lasted about an hour.

3.2 Experimental Hypotheses

We derive our experimental hypotheses by distilling Proposition 1 and the relevant lemma into comparative statics in the three choice variables in the game, the expert's probability

[19] Appendix B contains the experimental instructions sampled from treatment *NMP-SHUF-L*.

of choosing $r_{A\|B}$, s, the consumer's probability of acquiring the report, q, and her probability of choosing Product A, a. The relative frequencies of aggregate choices serve as the empirical counterparts of these probabilities, which we denote by s_T, q_T, and a_T, where $T = H, L, I$ abbreviates, respectively, *NMP-SHUF-H*, *NMP-SHUF-L*, and *NMP-INVA*.

Our first hypothesis compares, between the two main treatments, choices made at the strategically simultaneous stage. Despite the multiplicity of equilibria, not least anything chosen by the expert is consistent with equilibrium, the predicted relationship between the expert's and the consumer's choices serves to restrict observed behavior. The threshold in the expert's mixed strategy, $\bar{s}$, above which the consumer acquires the report, increases as $\bar{v}_B$ decreases. Our hypothesis leverages this relationship:

Hypothesis 1. *If the relative frequency of consumers' report acquisitions is at least as high in NMP-SHUF-H as in NMP-SHUF-L ($q_H \geq q_L$), then the relative frequency of experts' choices of $r_{A\|B}$ will be higher in NMP-SHUF-H ($s_H > s_L$).*

In deriving a hypothesis on noisy laboratory behavior from the precise theoretical statement, we posit that each consumer-subject has his/her own decision threshold that may deviate from and surround the theoretical values, giving rise to a distribution of thresholds in each treatment.[20] The point prediction translates into a directional prediction that the distribution of the thresholds in *NMP-SHUF-H* is a rightward shift of that in *NMP-SHUF-L*. While these thresholds are *unobservables*, the rationale behind Hypothesis 1 is that, if there are at least as many thresholds that fall on the side of acquiring the report in *NMP-SHUF-H* as in *NMP-SHUF-L*, reflected in the aggregate *observables* that $q_H \geq q_L$, there must be more choices of $r_{A\|B}$ by the experts in *NMP-SHUF-H* than in *NMP-SHUF-L*.[21]

A property of *NMP-INVA* that distinguishes it from *NMP-SHUF-H* and *NMP-SHUF-L* is that the consumer's default product choice can now be either product. In addition to the comparative statics with respect to $\bar{v}_B$, our second hypothesis, which compares *NMP-INVA* with each *NMP-SHUF-H* and *NMP-SHUF-L*, leverages this difference.

Hypothesis 2. *To compare NMP-INVA with the two main treatments,*

(a) if the relative frequency of consumers' report acquisitions is at least as high in NMP-SHUF-H/L as in NMP-INVA ($q_T \geq q_I$, $T = H, L$), then the relative frequency of experts' choices of $r_{A\|B}$ will be higher in NMP-SHUF-H/L ($s_T > s_I$, $T = H, L$); and

[20]For the adopted parameters, the theoretical values of the thresholds are $\bar{s}_H = \frac{35}{63}$ in *NMP-SHUF-H* and $\bar{s}_L = \frac{35}{69}$ in *NMP-SHUF-L*. With $\bar{s}_H > \bar{s}_L$, theory predicts that for $s \in (\bar{s}_L, \bar{s}_H)$ the consumer acquires the report in *NMP-SHUF-L* but not in *NMP-SHUF-H*.

[21]Under random matchings, each consumer effectively plays against a population of experts, which makes the experts' average behavior—the relative frequencies of ranking-method choices—relevant for the report-acquisition decisions made by each consumer.

(b) *conditional on not acquiring the report, the relative frequency of Product A being chosen by the consumers is lower in NMP-INVA than in NMP-SHUF-H/L* ($a_I(\varnothing) < a_T(\varnothing)$, $T = H, L$).

Hypothesis 2(b) is based on Lemma 2. Given the parameters of the treatments, the lemma predicts that Product A is always the consumer's default choice for *NMP-SHUF-H* and *NMP-SHUF-L*, while Product B could be the default choice for *NMP-INVA*. The consumer always acquires the report when Product B is the default. When Product A is the default, the same comparative statics behind Hypothesis 1 applies, where the threshold in the expert's mixed strategy to determine the consumer's report-acquisition decision, $\bar{s}$, is decreasing in $\bar{v}_B$, and this yields Hypothesis 2(a).[22]

4 Experimental Findings

We begin by examining aggregate behavior in Section 4.1, using the relative frequencies of aggregate choices, s_T, q_T, and a_T, $T = H, L, I$, to evaluate the experimental hypotheses. We then analyze individual behavior in Section 4.2, conducting a cluster analysis and estimating regressions to shed light on the constituents of the aggregate observations.

4.1 Aggregate Findings

Figure 1 presents the aggregate choices at the strategically simultaneous stage. We report our first finding, which addresses Hypotheses 1 and 2(a).

Finding 1. *For the strategically simultaneous stage, pairwise comparisons among the three treatments in respect of the relative frequencies of consumers' report acquisitions* (q_T) *and experts' choices of* $r_{A||B}$ (s_T) *are:*

(a) *NMP-SHUF-H vs. NMP-SHUF-L:* $q_H > q_L$ *and* $s_H > s_L$;

(b) *NMP-SHUF-L vs. NMP-INVA:* $q_I > q_L$ *and* $s_I > s_L$; *and*

(c) *NMP-SHUF-H vs. NMP-INVA:* $q_I > q_H$ *and* $s_I > s_H$.

[22]For *NMP-INVA*, there are three different thresholds in the expert's strategy, a lower bound $\bar{s}_I = \frac{35}{81}$, another lower bound for a different purpose $\tilde{s}_I = \frac{139}{150}$, and an upper bound $\hat{s}_I = \frac{104}{69}$ which is non-binding because it is larger than one. Skipping for brevity the cases of mixing, the following states how these thresholds operate: (a) for $s \in [0, \bar{s}_I)$, the consumer does not acquire the report under default Product A, (b) for $s \in (\bar{s}_I, \tilde{s}_I)$, she acquires the report under default Product A, and (c) for $s \in (\tilde{s}_I, 1]$, she acquires the report under default Product B. Insofar as the report-acquisition decision is concerned, $\tilde{s}_I$ is not relevant, and Hypothesis 2(a) follows the same rationale as Hypothesis 1 based on the fact that $\bar{s}_I < \bar{s}_L < \bar{s}_H$.

While only two comparisons, $s_H > s_L$ in (a) and $s_I > s_L$ in (b), are statistically significant, the corresponding combined comparisons are in line with Hypotheses 1 and 2(a).

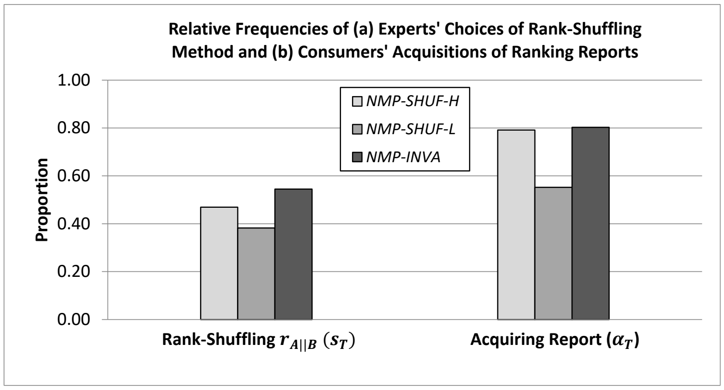

Figure 1: Aggregate Choices at Strategically Simultaneous Stage

Finding 1(a) compares the two main treatments. For consumers' report acquisitions, the relative frequency is 47% in *NMP-SHUF-H*, which is not significantly higher than the 38% in *NMP-SHUF-L* ($p = 0.2$, Mann-Whitney test). On the other hand, for experts' choices of $r_{A||B}$, the relative frequency is 79% in *NMP-SHUF-H*, significantly higher than the 55% in *NMP-SHUF-L* ($p = 0.05$, Mann-Whitney test). The joint comparisons, with the second comparison being statistically significant, are in line with Hypothesis 1: when a rank-shuffling method has a higher extent of misguidance, we observe more frequent choices of $r_{A||B}$ with no less frequent report acquisitions.

Findings 1(b) and 1(c) compare *NMP-INVA* with each of the main treatments. For report acquisitions, the relative frequency is 54% in *NMP-INVA*, which is not significantly higher than the 38% in *NMP-SHUF-L* and the 47% in *NMP-SHUF-H* ($p \geq 0.2$, Mann-Whitney tests). For choices of $r_{A||B}$, the relative frequency of 80% in *NMP-INVA* is significantly higher than the 55% in *NMP-SHUF-L* ($p = 0.05$, Mann-Whitney test) but not the 79% in *NMP-SHUF-H* ($p = 0.5$, Mann-Whitney test). The comparisons between *NMP-INVA* and *NMP-SHUF-L* lend some support to Hypothesis 2(a) in the form of the contrapositive of the conditional statement.[23] The comparisons between *NMP-INVA* and *NMP-SHUF-L* do not, however, support the hypothesis with statistical significance.

We turn to the stage where consumers choose products. Figure 2 presents the aggregate product choices. We report our second finding:

[23]For *NMP-INVA* and *NMP-SHUF-L*, Hypothesis 2(a) states that, if $q_L \geq q_I$, then $s_L > s_I$. Finding 1(b) is thus consistent with the contrapositive that, if $s_I \geq s_L$, then $q_I > q_L$. However, unlike the support for Hypothesis 1, in which the strict-inequality comparison obtains statistical significance, here the statistical significance is only with the weak-inequality comparison in the contrapositive. Any support for Hypothesis 2(a) is thus weaker than that for Hypothesis 1.

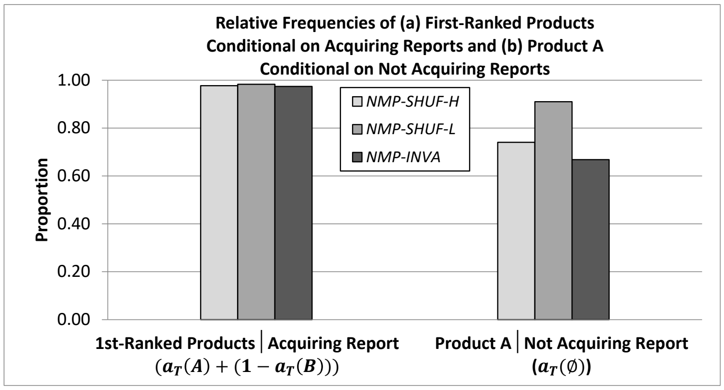

Figure 2: Aggregate Product Choices of Consumers

Finding 2. *For the product-choice stage,*

(a) conditional on the consumers acquiring ranking reports, the first-ranked products are nearly always chosen in the three treatments; and

(b) conditional on the consumers not acquiring ranking reports, the relative frequencies of Product A being chosen in the three treatments are in the order $a_L(\emptyset) > a_H(\emptyset) > a_I(\emptyset)$*, where only the difference between NMP-SHUF-L and NMP-INVA is statistically significant, supporting part of Hypothesis 2(b).*

The relative frequency of Product A being chosen is 67% in *NMP-INVA*, which is significantly lower than the 91% in *NMP-SHUF-L* ($p = 0.05$, Mann-Whitney test) but not the 74% in *NMP-SHUF-H* ($p = 0.2$, Mann-Whitney test). Finding 2(b) supports Hypothesis 2(b) in respect of the comparison between *NMP-SHUF-L* and *NMP-INVA*. While it is not a focus of our inquiry as there is no related comparative statics across treatments, Finding 2(a) shows that ranking reports influence consumers' product choices. Upon viewing the ranking reports, consumers choose the first-ranked products 98% of the time in the two main treatments and 97% of the time in *NMP-INVA*. Report-acquisition decisions are apparently motivated by product choices.

We conclude our analysis of aggregate behavior by interpreting findings from the two main treatments in terms of the phenomenon of shuffling to sell in the game. Theoretically given the treatment parameters, the consumer, when not acquiring the report, is worse off by 150 expected payoff-points under $r_{A\|B}$ than under r_A in both treatments; when acquiring the report, she is worse off by 24 and 12 in *NMP-SHUF-H* and *NMP-SHUF-L* respectively. By choosing $r_{A\|B}$ instead of r_A, the expert effectively creates two unfavorable situations (being worse off by 150 vs. 24 or 12) and offers the less unfavorable with a price of 70.

Our findings from *NMP-SHUF-H* and *NMP-SHUF-L* indicate that this shuffling to sell occurs in the laboratory. More importantly, the comparative statics in Finding 1(a) says that the experts in *NMP-SHUF-H* shuffle more often, when doing so hurts the consumers more. The following finding reports the payoff consequences of this observation:

Finding 3. *Consumers' average payoffs are lower in NMP-SHUF-H than in NMP-SHUF-L, while the opposite is observed for experts' average payoffs.*

For consumers, the average payoffs are 257.85 in *NMP-SHUF-L* and 235.20 in *NMP-SHUF-H*. Even if we adjust for the lower parameter value of $\bar{v}_B$ in *NMP-SHUF-H*, the consumers in this treatment still score a lower average payoff of 241.02—they are made worse off *endogenously*.[24] At the same time, experts' average payoffs are higher in *NMP-SHUF-H* than in *NMP-SHUF-L* (164.22 vs. 133.73). This comparative statics exemplifies the benefit the rank shuffling brings to experts at the expense of consumers.

4.2 Individual Analysis

We proceed to examine the constituents of the aggregate observations by analyzing subject-level data. We begin with a visual categorization of individual subjects using the algorithm of k-means clustering (MacQueen, 1967). We then perform a regression analysis to further examine quantitatively the patterns of individual behavior.

k-Means Clustering. We compute the relative frequency of each subject's choices across all rounds and use it as a proxy of either, depending on the subject's role, the expert's mixed strategy or the relevant part of the consumer's behavior strategy. The observations, one for each subject, are partitioned into a pre-determined number (k) of clusters based on the proximity of each observation to the center (mean) of a cluster. We find that setting $k = 3$ categorizes our data reasonably well. Each decision in our game is binary with two pure choices; using three clusters also has the intuitive appeal that we can visualize which cluster is closest to one of the pure choices and which one is in between.

Figure 3 presents the consumer clusters in three panels, one for each treatment, where we plot the relative frequency of Product A chosen as defaults against the relative frequency of report acquisitions.[25] Each marker in the figures represents a subject. The grayed markers are cluster means. We refer to the *size* of a cluster by the number of subjects it contains (e.g., the cluster with the most subjects is of the largest size).

The k-means clustering is performed based on the frequency *pair* of each consumer. For

[24] The "counter-factual" average payoff is obtained by recomputing the consumers' average payoff in *NMP-SHUF-H* using 80, the value of $\bar{v}_B$ in *NMP-SHUF-L*, instead of the original 60.

[25] Given that the first-ranked products are nearly always chosen conditional on the acquisitions of reports [Finding 2(a)], we exclude this part of the strategy in our analysis; the minimal variations do not afford a meaningful differentiation of behavior.

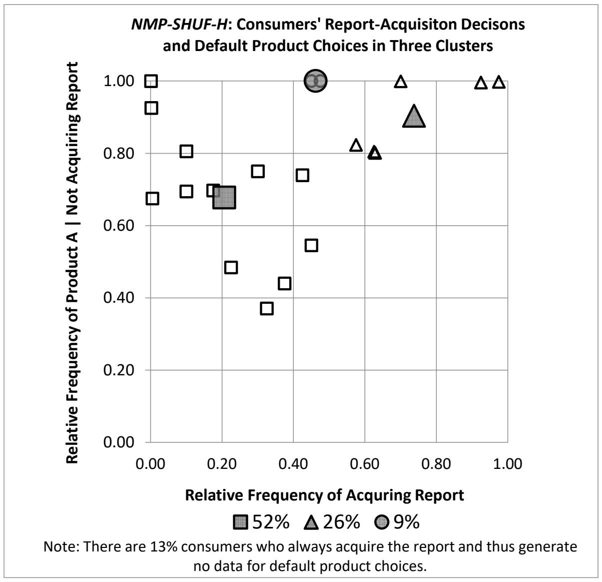

(a) *NMP-SHUF-H*

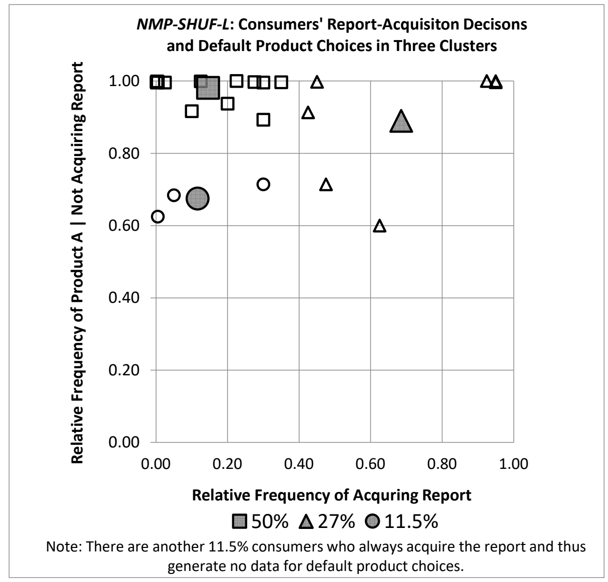

(b) *NMP-SHUF-L*

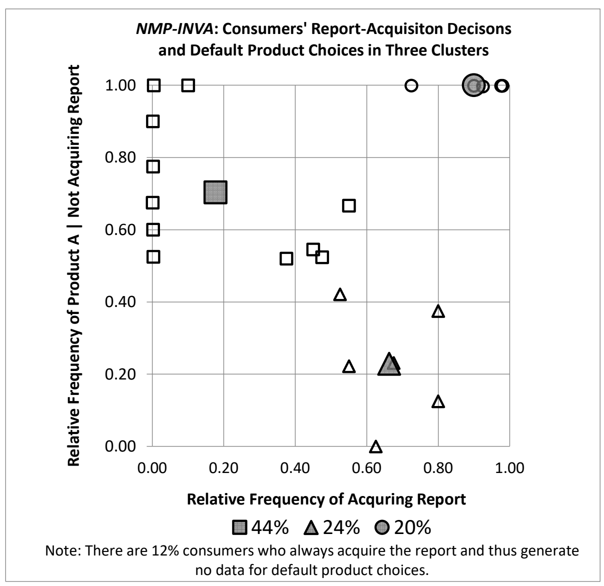

(c) *NMP-INVA*

Figure 3: Three-Means Clustering: Consumers

expositional purpose, however, we present our observations separately for report acquisitions and default product choices. We start with the former:

Finding 4. *For ranking-report acquisitions, individual decisions vary substantially within treatments but are similar across treatments. The largest cluster in each of the three treatments commonly contains consumers who infrequently acquire ranking reports, while the second largest cluster commonly contains consumers who often acquire reports.*

Finding 4 indicates that the aggregate observations on report acquisitions presented in Section 4.1 are made up of heterogenous individual behavior, which can largely be categorized into two groups with frequent and infrequent acquisitions. The heterogeneity is, however, similar across treatments. The modal individual behavior in each treatment, as captured by the respective largest clusters (marked by □ in Figure 3), features infrequent report acquisitions; the average acquisition frequencies in these □-clusters, which account for 52% of the consumers in *NMP-SHUF-H*, 50% in *NMP-SHUF-L*, and 44% in *NMP-INVA*, are 21%, 15%, and 18% respectively.

Frequent report acquisitions are observed among the second largest clusters (marked by △); the average acquisition frequencies in these △-clusters, which account for 26% of the consumers in *NMP-SHUF-H*, 27% in *NMP-SHUF-L*, and 24% in *NMP-INVA*, are 74%, 69%, and 66% respectively. In addition to these consumers who display the second most prevalent individual behavior as frequent report acquirers, there are also about 12% of consumers in each treatment who always acquire reports and thus are not amenable to the clustering exercise because of the missing default-product observations.[26] Overall, the within-treatment heterogeneity in report acquisitions is consistent with the thesis behind our hypotheses on aggregate behavior that there is among consumers a distribution of different decision thresholds.

We present next the observation on product choices when reports are not acquired:

Finding 5. *For default-product choices, individual decisions vary both within and across treatments. In terms of the largest clusters, NMP-SHUF-L is differentiated from the other treatments, in which the consumers therein nearly always choose Product A. In terms of the second largest clusters, NMP-INVA is differentiated from the other treatments, in which the consumers therein choose Product A substantially less often.*

Recall that the treatment variable, $\bar{v}_B$, is monotonically increasing from *NMP-SHUF-H* to *NMP-SHUF-L* to *NMP-INVA*. Equilibrium predictions aside, one may expect that an increase in the expected monetary reward from choosing Product B may lead to less frequent choices of Product A. Among the largest clusters, however, the differentiated observation from *NMP-SHUF-L* presents a "non-monotonicity" in this regard; the average frequencies of Product A being chosen in these □-clusters are 68% in *NMP-SHUF-H*, 98% in *NMP-SHUF-L*, and 70% in *NMP-INVA*.

While this non-monotonicity in modal behavior cannot be explained by naive incentive pursuits or equilibrium, observations from the second largest clusters are consistent with a

[26]The behavior in the smallest clusters (marked by ○), which comprise 9% of the consumers in *NMP-SHUF-H*, 12% in *NMP-SHUF-L*, and 20% in *NMP-INVA*, has more variations across treatments. The average acquisition frequencies in these ○-clusters are 46% in *NMP-SHUF-H*, 12% in *NMP-SHUF-L*, and 90% in *NMP-INVA*.

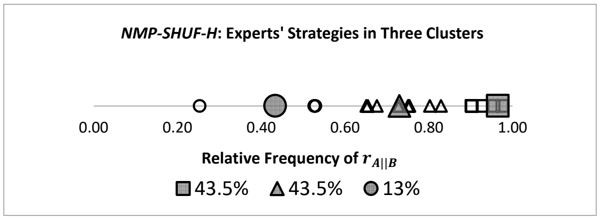

(a) *NMP-SHUF-H*

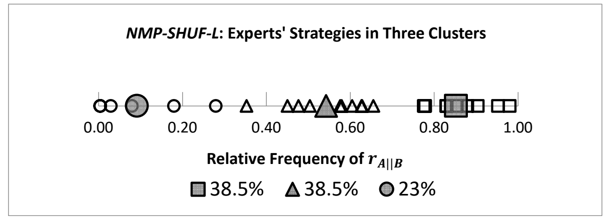

(b) *NMP-SHUF-L*

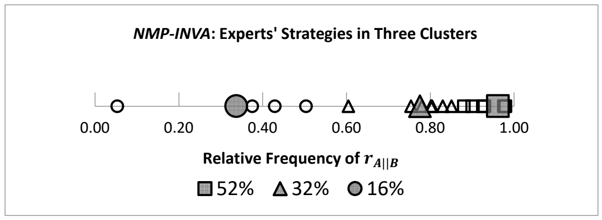

(c) *NMP-INVA*

Figure 4: Three-Means Clustering: Experts

qualitative property of equilibrium product choices. Theoretically, Product A is the only default choice in *NMP-SHUF-H* and *NMP-SHUF-L*, while Product B might also be chosen as a default in *NMP-INVA*; empirically among the △-clusters, the average frequencies of Product A being chosen are—despite the variation in $\bar{v}_B$—almost identical in *NMP-SHUF-H* (90%) and *NMP-SHUF-L* (89%), while it is distinctively lower in *NMP-INVA* (23%).[27]

Turning to experts, Figure 4 presents the simple clustering outcomes based on the relative frequencies of the rank-shuffling $r_{A||B}$ being chosen by individual experts in all rounds. We summarize our observation:

Finding 6. *For experts' ranking-method choices,*

(a) individual decisions vary within treatments, in which the cluster-average frequency of the ranking-shuffling $r_{A||B}$ decreases in cluster size in each treatment; sorting the frequencies across all clusters of all treatments, the most frequent choices of $r_{A||B}$ are recorded in the largest cluster in NMP-SHUF-H and the least frequent recorded in the smallest cluster in NMP-SHUF-L; and

(b) individual decisions vary to a lesser extent across treatments, in which, for a given rank in size, the clusters in NMP-SHUF-H and NMP-INVA share similar average

[27] *NMP-SHUF-L* is differentiated also with respect to the smallest clusters (○), this time by containing consumers who less often choose Product A as the default. The average frequencies of Product A being chosen in these ○-clusters are 100% in *NMP-SHUF-H*, 67% in *NMP-SHUF-L*, and 100% in *NMP-INVA*. They are, however, based on small numbers of observations; the ○-clusters in *NMP-SHUF-H*, *NMP-SHUF-L*, and *NMP-INVA* contain, respectively, 9%, 12%, and 20% of the consumers (despite the relatively higher proportion in *NMP-INVA*, those consumers on average make a default choice only 10% of the time).

frequencies of $r_{A\|B}$, which are in turn higher than the average frequency of $r_{A\|B}$ in the corresponding clusters in NMP-SHUF-L.

Although there is only one choice variable by the experts, the k-means clustering, ostensibly an overkill, helps dissect the aggregate findings in a way that otherwise may not be apparent. For the two main treatments, the categorizations of choices allow us to distinguish between two observations that together fuel the comparison in Finding 1(a): the more frequent aggregate choices of $r_{A\|B}$ in *NMP-SHUF-H* than in *NMP-SHUF-L* are the result of proportionally more experts in the former in the larger clusters, i.e., more experts who frequently choose $r_{A\|B}$ [Finding 6(a)], *and* they choose $r_{A\|B}$ with higher average frequencies than their counterparts in *NMP-SHUF-L* [Finding 6(b)]. In *NMP-SHUF-H*, there are 43.5% of experts in each of the two largest clusters (marked by □ and △ in Figure 4) and 13% in the smallest (○), and the average frequencies of $r_{A\|B}$ are 97% in the □-cluster, 73% in the △-cluster, and 43% in the ○-cluster. On the other hand, in *NMP-SHUF-L*, there are 38.5% of experts in each of the two largest clusters and 23% in the smallest, and the cluster-average frequencies of $r_{A\|B}$ are, respectively, 85%, 54%, and 9%.

The clustering also sheds light on the aggregate comparisons between *NMP-INVA* and the main treatments. In *NMP-INVA*, there are 52% of experts in the largest cluster (□), 32% in the second largest (△), and 16% in the smallest (○). The average frequencies of $r_{A\|B}$ are 96% in the □-cluster, 78% in the △-cluster, and 34% in the ○-cluster. The profile of these clusters is similar to that in *NMP-SHUF-H*, and this accounts for the insignificant difference in the aggregate frequencies of $r_{A\|B}$ between the two treatments. Given this similarity, the observations above regarding *NMP-SHUF-H* and *NMP-SHUF-L* would also apply to dissect the aggregate differences between *NMP-INVA* and *NMP-SHUF-L*: relative to *NMP-SHUF-L*, there are proportionally more experts in *NMP-INVA* who choose $r_{A\|B}$ frequently, and they do so with higher frequencies.

Regression Analysis. We further analyze subject-level data with regressions, complementing the visual categorization with a quantitative analysis. We estimate the following linear probability model using the panel choice data made by matched groups in 40 rounds:

$$S_{i\tau} = \alpha_0 + \alpha_1 Q_{i\tau} + \alpha_2 \mathcal{H}_i + \alpha_3 \mathcal{I}_i + \alpha_5 (Q_{i\tau} \times \mathcal{H}_i) + \alpha_6 (Q_{i\tau} \times \mathcal{I}_i) + \varepsilon_{i\tau}. \tag{1}$$

Except for the error term, all the variables in equation (1) are dummy variables: $S_{i\tau}$ takes the value of one if expert i chooses $r_{A\|B}$ ($s = 1$) in round τ, $Q_{i\tau}$ takes the value of one if the consumer matched with expert i in round τ acquires the ranking report ($q = 1$), $\mathcal{H}_i$ takes the value of one if expert i is in treatment *NMP-SHUF-H*, and $\mathcal{I}_i$ takes the value of one if expert i is in treatment *NMP-INVA*.

Table 4: Estimation Results of Equation (1)

	(1)	(2)	(3)
Constant	0.555***	0.556***	–
	(0.017)	(0.064)	–
$Q_{i\tau}$	0.007	0.004	0.004
	(0.028)	(0.032)	(0.028)
$\mathcal{H}_{i\tau}$	0.223***	0.238**	0.233***
	(0.026)	(0.079)	(0.061)
$\mathcal{I}_{i\tau}$	0.221***	0.237**	0.239***
	(0.027)	(0.084)	(0.072)
$Q_{i\tau} \times \mathcal{H}_{i\tau}$	0.025	−0.006	−0.003
	(0.040)	(0.034)	(0.030)
$Q_{i\tau} \times \mathcal{I}_{i\tau}$	0.042	0.013	0.020
	(0.039)	(0.037)	(0.034)
No. of Observations	2,960	2,960	2,960

Note: The dependent variable is a dummy variable for the choice of $r_{A\|B}$. Column (1) reports the OLS estimates. Column (2) reports the random-effects estimates. Column (3) reports the average marginal effects from estimating the random-effects probit model (no constant term is included because the reported numbers are marginal effects). Robust standard errors clustered at the session level are in parentheses. *** indicates significance at 0.1% level, ** significance at 1% level, and * significance at 5% level.

The specification of the regression is motivated by theory and the aggregate findings. There are significant differences between *NMP-SHUF-L* and each of *NMP-SHUF-H* and *NMP-INVA* (but not between *NMP-SHUF-H* and *NMP-INVA*) in the aggregate choices of ranking methods. Accordingly, the regression uses *NMP-SHUF-L* as the baseline and evaluates it against *NMP-SHUF-H* and *NMP-INVA*.[28] As far as the theoretical predictions are concerned, we are in essence reevaluating the experimental hypotheses using the panel data. Theory suggests the comparative statics that, controlling for $Q_{i\tau}$, the probability that the expert chooses $r_{A\|B}$ should be higher in *NMP-SHUF-H* than in *NMP-SHUF-L*; α_2 is expected to be positive. The same comparative statics would predict α_3 to be negative, although the aggregate finding, which does not control for the consumers' decisions, suggests that it might be otherwise. The equilibrium multiplicity does not allow us to pin down the sign of α_1. Theory is also silent on the effects of the interaction terms.

Table 4 reports the estimation results under three different estimation methods, OLS, random-effects, and the probit (random-effects) counterpart of the linear probability model. All three methods yield similar results. The sign of the estimated α_2, which is positive with statistical significance, is consistent with the comparative-statics prediction and the aggregate finding. On the other hand, the statistically significant estimate of α_3, which is also positive, is not as theoretically predicted. It is, however, in line with the aggregate finding. In Section 4.1, we document that the aggregate frequency of $r_{A\|B}$ is higher in *NMP-*

[28]The constant term α_0 measures the probability of $r_{A\|B}$ being chosen in *NMP-SHUF-L* when the consumer does not acquire the ranking report.

INVA than in *NMP-SHUF-L* under a higher aggregate frequency of report acquisitions in the former. The fact that the estimated α_3 is positive indicates that the analogous estimated probability of $r_{A||B}$ is higher in *NMP-INVA* even when we control for consumers' report-acquisition decisions. The already weak support for Hypothesis 2(a) in the form of its contrapositive vanishes at the individual level. We conclude our data analysis with the following summary of the regression findings:

Finding 7. *The regression analysis corroborates the aggregate findings that the experts in NMP-SHUF-H and NMP-INVA choose the rank-shuffling $r_{A||B}$ more frequently than those in NMP-SHUF-L. It provides further support for Hypothesis 1 at the subject level in respect of the comparative statics between the two main treatments.*

5 Concluding Remarks

Motivated by the lack of structured evidence on the sentiment expressed by some commentators that ranking publishers excessively alter their product rankings for marketing purposes, this study resorts to laboratory evidence. We use monetary payments to induce in the laboratory plausible incentives faced by ranking publishers. The experimental design is guided by the formal analysis of a ranking-report game, which helps make precise the layman view. The expert in our game "alters the product rankings" in the manner of choosing a rank-shuffling method sufficiently often. This equilibrium strategy, under which the expert sells his ranking report, induces the consumer to pay to resolve an uncertainty over a ranking value interpreted as prestige. The altering is done "excessively" in the sense that it involves misguiding the consumer, who is worse off relative to the case where the expert adopts a method that does not shuffle yet always provides product guidance.

Despite this distinct message from the theoretical analysis, there is an immense equilibrium multiplicity in our game. As such, in deriving and evaluating experimental hypotheses, our focus has been on the comparative statics rather than point predictions. Equilibria predict that, as the rank-shuffling method serves less of the guidance purpose, selling the ranking report calls for more frequent adoptions of the method, and consumers would be more worse off as a result. We observe these two comparative statics in our two main treatments. Experts in the treatment where the rank-shuffling method is associated with a higher degree of misguidance choose the method more frequently. This behavior of the experts causes the consumers to be worse off endogenously relative to the consumers in the other treatment with a less misguiding rank-shuffling method. Findings from our third, exploratory treatment, obtained under treatment parameters that allow for an even wider scope of equilibrium multiplicity, are nevertheless less in line with the predictions.

Overall, our experiment provides evidence supporting the view that a profit-driven ranking publisher may adopt a ranking methodology to facilitate sales at the expense of consumers.

Our findings are obtained and interpreted under the posited existence of a channel distinct from product information—the ranking value interpreted as prestige—via which product rankings influence consumers, and, underscoring the key insight of our theoretical and empirical analysis, this creates an incentive for the expert to inject uncertainty into the rankings. It is worth noting that we have not provided evidence for the existence of such a prestige effect; rather, we have examined the possible consequences of it when it exists. There may be other channels through which product rankings influence consumer choices. In the area of university rankings, e.g., Luca and Smith (2013) obtain evidence supporting the view that the salience provided by a simple rank order plays a role in influencing students' decisions. Interestingly, this same evidence is also consistent with the view that students' decisions are influenced by the prestige associated with attending universities ranked top by popular university-ranking publications. To fully understand the incentives and behavior surrounding product rankings, one direction of future empirical research, either inside or outside of the laboratory, is to comprehensively identify the channels of influences of product rankings on consumer choices.

In our theory and experiment, the sellers of the products being ranked are not part of the model. Another direction of future research is to study a richer environment in which sellers are players themselves. Product sellers may play a strategic role in the impacts of product rankings on consumers. Luca and Smith (2015), e.g., provide empirical evidence that business schools selectively promote the publications in which their MBA programs are favorably ranked. Questions of interest that are relevant to our study, which could be addressed theoretically and experimentally, include how sellers strategically respond when ranking publishers engage in shuffling, whether those responses strengthen or attenuate the incentives to shuffle, and the resulting welfare implications for consumers.

References

Anderson, Eric T. and Duncan I. Simester (2014), "Reviews without a purchase: Low ratings, loyal customers, and deception." *Journal of Marketing Research*, 51, 249–269.

Au, Pak Hung and King King Li (2018), "Bayesian persuasion and reciprocity: Theory and experiment." Working Paper.

Bagwell, Kyle (2007), "The economic analysis of advertising." In *Handbook of Industrial Organization*, volume 3, 1701–1844, Elsevier.

Banks, Jeffrey C., Colin F. Camerer, and David Porter (1994), "Experimental tests of nash refinements in signaling games." *Games and Economic Behavior*, 6, 1–31.

Becker, Gary S. and Kevin M. Murphy (1993), "A simple theory of advertising as a good or bad." *The Quarterly Journal of Economics*, 108, 941–964.

Blume, Andreas, Oliver J. Board, and Kohei Kawamura (2007), "Noisy talk." *Theoretical Economics*, 2, 395–440.

Blume, Andreas, Ernest K. Lai, and Wooyoung Lim (2020), "Mediated talk: An experiment." Working Paper.

Brandts, Jordi and Charles A. Holt (1992), "An experimental test of equilibrium dominance in signaling games." *American Economic Review*, 82, 1350–1365.

Chen, Daniel L., Martin Schonger, and Chris Wickens (2016), "oTree—An open-source platform for laboratory, online, and field experiments." *Journal of Behavioral and Experimental Finance*, 9, 88–97.

Chevalier, Judith A and Dina Mayzlin (2006), "The effect of word of mouth on sales: Online book reviews." *Journal of Marketing Research*, 43, 345–354.

Crawford, Vincent P. and Joel Sobel (1982), "Strategic information transmission." *Econometrica*, 1431–1451.

Dearden, James A., Rajdeep Grewal, and Gary L. Lilien (2019), "Strategic manipulation of university rankings, the prestige effect, and student university choice." *Journal of Marketing Research*, 56, 691–707.

Dickhaut, John W., Kevin A. McCabe, and Arijit Mukherji (1995), "An experimental study of strategic information transmission." *Economic Theory*, 6, 389–403.

Forsythe, Robert R., Mark Isaac, and Thomas R. Palfrey (1989), "Theories and tests of 'blind bidding' in sealed-bid auctions." *RAND Journal of Economics*, 20, 214–238.

Fréchette, Guillaume, Alessandro Lizzeri, and Jacopo Perego (2020), "Rules and commitment in communication: An experimental analysis." Working Paper.

Gneezy, Uri (2005), "Deception: The role of consequences." *American Economic Review*, 95, 384–394.

Goltsman, Maria, Johannes Hörner, Gregory Pavlov, and Francesco Squintani (2009), "Mediation, arbitration, and negotiation." *Journal of Economic Theory*, 144, 1397–1420.

Grossman, Sanford J. (1981), "The informational role of warranties and private disclosure about product quality." *Journal of Law & Economics*, 24, 461–483.

Kamenica, Emir and Matthew Gentzkow (2011), "Bayesian persuasion." *American Economic Review*, 101, 2590–2615.

King, Ronald R. and David E. Wallin (1991), "Voluntary disclosures when seller's level of information is unknown." *Journal of Accounting Research*, 29, 96–108.

Krishna, Vijay and John Morgan (2004), "The art of conversation: Eliciting information from experts through multi-stage communication." *Journal of Economic Theory*, 1417, 147–179.

Kuksov, Dmitri and Kangkang Wang (2013), "A model of the "it" products in fashion." *Marketing Science*, 32, 51–69.

Kuksov, Dmitri and Ying Xie (2012), "Competition in a status goods market." *Journal of Marketing Research*, 49, 609–623.

Lu, Susan F. and Huaxia Rui (2018), "Can we trust online physician ratings? Evidence from cardiac surgeons in Florida." *Management Science*, 64, 2557–2573.

Luca, Michael and Jonathan Smith (2013), "Salience in quality disclosure: Evidence from the US News college rankings." *Journal of Economics & Management Strategy*, 22, 58–77.

Luca, Michael and Jonathan Smith (2015), "Strategic disclosure: The case of business school rankings." *Journal of Economic Behavior & Organization*, 112, 17–25.

Luca, Michael and Georgios Zervas (2016), "Fake it till you make it: Reputation, competition, and Yelp review fraud." *Management Science*, 62, 3412–3427.

MacQueen, James (1967), "Some methods for classification and analysis of multivariate observations." *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, 1, 281–297.

Mayzlin, Dina, Yaniv Dover, and Judith Chevalier (2014), "Promotional reviews: An empirical investigation of online review manipulation." *American Economic Review*, 104, 2421–55.

Milgrom, Paul R. (1981), "Good news and bad news: Representation theorems and applications." *Bell Journal of Economics*, 12, 380–391.

Nguyen, Quyen (2016), "Bayesian persuasion: Evidence from the laboratory." Working Paper.

Pope, Devin G. (2009), "Reacting to rankings: Evidence from 'America's best hospitals'." *Journal of Health Economics*, 28, 1154–1165.

Sahoo, Nachiketa, Chrysanthos Dellarocas, and Shuba Srinivasan (2018), "The impact of online product reviews on product returns." *Information Systems Research*, 29, 723–738.

Spence, Michael (1973), "Job market signaling." *Quarterly Journal of Economics*, 87, 355–374.

Sun, Monic (2012), "How does the variance of product ratings matter?" *Management Science*, 58, 696–707.

Tierney, John (2013), "Your annual reminder to ignore the US News & World Report college rankings." *The Atlantic*, URL https://www.theatlantic.com/education/archive/2013/09/your-annual-reminder-to-ignore-the-em-us-news-world-report-em-college-rankings/279103/. Last accessed January 5, 2021.

Zhu, Feng and Xiaoquan Zhang (2010), "Impact of online consumer reviews on sales: The moderating role of product and consumer characteristics." *Journal of Marketing*, 74, 133–148.

Appendix A Proofs

Proof of Lemma 1. Consider first the case where the consumer has viewed report A. Bayes' rule implies that $\mu_{\bar{v}_B} = \frac{p-ps}{1-ps}$, and the consumer's expected payoffs from choosing Products A and B are, respectively, $\bar{v}_A + \gamma - c$ and $\left(\frac{p-ps}{1-ps}\right)\bar{v}_B - c$. Given that $\gamma > 0$ and $\bar{v}_A > p\bar{v}_B$, for all $s \in [0,1]$, the former payoff is strictly higher than the latter. Thus, $a(A) = 1$ is the unique best response for any $s \in [0,1]$.

Consider next the case where the consumer has viewed report B. Bayes' rule implies that $\mu_{\bar{v}_B} = 1$ for $s > 0$. For $s = 0$, since the off-path information set is a singleton, the product belief must also be that $\mu_{\bar{v}_B} = 1$. The consumer's expected payoffs from choosing Products A and B are, respectively, $\bar{v}_A - c$ and $\bar{v}_B + \gamma - c$. It follows that (a) if $\gamma > \bar{v}_A - \bar{v}_B$, the consumer's unique best response is $a(B) = 0$; (b) if $\gamma < \bar{v}_A - \bar{v}_B$, her unique best response is $a(B) = 1$; and (c) if $\gamma = \bar{v}_A - \bar{v}_B$, randomizing between the two products with any probability, $a(B) \in [0,1]$, constitutes a best response.

The three cases in the lemma obtain by pairing the case of report A with each of the three cases of report B.

□

Proof of Lemma 2. Without viewing the report, the consumer's product belief $\mu_{\bar{v}_B} = p$. Note also that her strategy belief $\mu_{r_{A\|B}} = s$. Given these beliefs, her expected payoffs from purchasing Products A and B are, respectively, $\bar{v}_A + [s(1-p) + (1-s)]\gamma$ and $p(\bar{v}_B + s\gamma)$. It follows that, for $\tilde{s} = \frac{\bar{v}_A - p\bar{v}_B + \gamma}{2p\gamma} > 0$, (a) if $s < \tilde{s}$, the consumer strictly prefers Product A; (b) if $s > \tilde{s}$, she strictly prefers Product B; and (c) if $s = \tilde{s}$, she is indifferent between the two products. Accordingly, (a) if $\bar{v}_A - p\bar{v}_B > (2p-1)\gamma$ so that $\tilde{s} > 1$, $a(\varnothing) = 1$ is the consumer's unique best response for all $s \in [0,1]$; and (b) if $\bar{v}_A - p\bar{v}_B \leq (2p-1)\gamma$ so that $\tilde{s} \leq 1$, (i) for $s \in [0, \tilde{s})$, the consumer's unique best response is $a(\varnothing) = 1$, (ii) for $s = \tilde{s}$, randomizing between the two products with any probability, $a(\varnothing) \in [0,1]$, constitutes a best response, and, (iii) for $s \in (\tilde{s}, 1]$ in the case where $\tilde{s} < 1$, her unique best response is $a(\varnothing) = 0$.

□

Proof of Proposition 1. For $a(\varnothing) = a(A) = 1$ and $a(B) = 0$, the consumer's expected payoffs from choosing $q = 1$ and $q = 0$ are, respectively, $s[p\bar{v}_B + (1-p)\bar{v}_A] + (1-s)\bar{v}_A + \gamma - c$ and $\bar{v}_A + [s(1-p) + (1-s)]\gamma$. It follows that, for $\bar{s} = \frac{c}{p(\bar{v}_B + \gamma - \bar{v}_A)}$, (a) if $s < \bar{s}$, she strictly prefers to choose $q = 0$; (b) if $s > \bar{s}$, she strictly prefers to choose $q = 1$; and (c) if $s = \bar{s}$, she is indifferent between acquiring and not acquiring the report. From Lemmas 1 and 2 and their proofs, her product-purchasing rule features (a) $a(A) = 1$ and $a(B) = 0$ as unique best response if and only if $\gamma > \bar{v}_A - \bar{v}_B$, and (b) $a(\varnothing) = 1$ as unique best response if and only if (i) $\bar{v}_A - p\bar{v}_B > (2p-1)\gamma$ or (ii) $\bar{v}_A - p\bar{v}_B \leq (2p-1)\gamma$ and $s < \tilde{s} = \frac{\bar{v}_A - p\bar{v}_B + \gamma}{2p\gamma}$. The condition that $\gamma > \bar{v}_A - \bar{v}_B$ implies that $\bar{q} > 0$, while that $\bar{v}_A - p\bar{v}_B > (2p-1)\gamma$ or $\bar{v}_A - p\bar{v}_B \leq (2p-1)\gamma$ does not by itself restrict the range of $\bar{s}$. Furthermore, we can have that $\bar{s} \gtreqless \tilde{s}$, and thus the parameter profiles for the product-purchasing rule are compatible with those for the report-acquisition decisions.

For $a(A) = 1$ and $a(\varnothing) = a(B) = 0$, the consumer's expected payoffs from choosing $q = 1$ and $q = 0$ are, respectively, $s[p\bar{v}_B + (1-p)\bar{v}_A] + (1-s)\bar{v}_A + \gamma - c$ and $p(\bar{v}_B + s\gamma)$. It follows that, for $\hat{s} = \frac{\bar{v}_A - p\bar{v}_B + \gamma - c}{p(\bar{v}_A + \gamma - \bar{v}_B)}$, (a) if $s < \hat{s}$, she strictly prefers to choose $q = 0$ ($q = 1$) in the case where $\bar{v}_A + \gamma < \bar{v}_B$ ($\bar{v}_A + \gamma \geq \bar{v}_B$); (b) if $s > \hat{s}$, she strictly prefers to choose $q = 1$ ($q = 0$) in the case where $\bar{v}_A + \gamma < \bar{v}_B$ ($\bar{v}_A + \gamma \geq \bar{v}_B$); and (c) if $s = \hat{s}$, she is indifferent between acquiring and not acquiring the report. From Lemma 2 and its proof, her product-purchasing rule features $a(\varnothing) = 0$ as unique best response if and only if $\bar{v}_A - p\bar{v}_B \leq (2p-1)\gamma$ and $s > \tilde{s} = \frac{\bar{v}_A - p\bar{v}_B + \gamma}{2p\gamma}$. The conditions that $\bar{v}_A - p\bar{v}_B \leq (2p-1)\gamma$ and, for $a(A) = 1$ and $a(B) = 0$, that $\gamma > \bar{v}_A - \bar{v}_B$ do not by themselves restrict the range of $\hat{s}$. Furthermore, we can have that $\hat{s} \gtreqless \tilde{s}$, and thus the parameter profiles for the product-purchasing rule and for the report-acquisition decisions are compatible. It is straightforward that the compatibility extends to the case where the consumer randomizes between the products when not viewing the report.

For the expert's best-responding choices of ranking method, note that, for a given report-acquisition decision of the consumer, the expert is indifferent between r_A and $r_{A\|B}$. Thus, randomizing between the two methods with any probability is a best response.

$\square$

Appendix B Sample Experimental Instructions: Treatment *NMP-SHUF-L*

INSTRUCTIONS

Welcome to this experiment, which studies decision making between two individuals. The experiment will last approximately 1.5 hour. There will be 40 rounds of decisions. Please read these instructions carefully. The cash payment you receive at the end of the experiment depends on your decisions.

Your Role and Decision Group

Half of the participants in today's session will be randomly assigned the role of **Product Expert**, and the other half the role of **Consumer**. Your role will remain fixed throughout the experiment. In each and every round, one Product Expert is matched with one Consumer. Participants will be randomly rematched after each round. You will not learn the identity of the participant you are matched with in any round, nor will that participant learn your identity.

Your Decision in Each Round

Overview. There are two products, A and B. They differ by the values to the Consumer. The value of Product A is fixed at 100.

The value of Product B is either 0 or 80. In each and every round, the computer will randomly select one of these two values for Product B according to the following:

(a) The chance that 0 will be selected is 40%.

(b) The chance that 80 will be selected is 60%.

The Product Expert makes one decision: how to rank the two products. The Consumer makes two decisions: whether to purchase to view the Product Expert's ranking report and which product to choose.

Product Expert's Decision: Choosing a Ranking Method. In each and every round, you choose a **ranking method**. There are two choices:

(a) Method 1: Rank Product A first regardless of the value of Product B.

(b) Method 2:

  (i) Rank Product A first if the value of Product B is selected to be 0.

  (ii) Rank Product B first if the value of Product B is selected to be 80.

Your decision therefore comes down to which product to rank first in the case that the value of Product B is selected to be 80.

You choose a ranking method before the computer selects the value for Product B, i.e., you don't know the selected value of Product B when you choose. Your task for the round is completed after you make the choice.

Depending on the selected value of Product B, a **ranking report** is generated according to your chosen ranking method. If the Consumer chooses to purchase to view your ranking report, the generated report, either "Product A is ranked 1st" or "Product B is ranked 1st," will be revealed to the Consumer.

Note that during the round the Consumer will never see the selected value of Product B. Note also that even if the Consumer chooses not to purchase your ranking report, the product ranking will still affect him/her. This will be further explained below.

Consumer's Decisions: Whether to Purchase the Ranking Report and Choosing a Product. In each and every round, your first decision is whether to purchase to view the ranking report (Purchase or Pass). You need to pay for the report if you choose "Purchase," which will be further explained below.

Irrespective of your first decision, you will make a second decision on which product, A or B, to choose. While you will receive the value of the product you choose, both products are free—you don't need to pay for it.

If you purchase the ranking report, you will make your product choice after seeing the generated report (you will only see the ranking report, not the ranking method that generates it). If you pass, you will make your product choice right after you choose "Pass." In either case, you will make your product choice without seeing directly the selected value of Product B. Your task for the round is completed after you make your product choice.

Your Reward in Each Round

Your reward in each round is expressed in "experimental currency unit" (ECU). How your earned ECU converts into cash payment will be explained below.

Product Expert's Reward. You will receive the amount the Consumer pays for the report, 70 ECU, multiplied by 5. Thus, if the Consumer chooses to purchase your ranking report, you will receive a total of 350 ECU for the round.

Table B.1 summarizes your potential reward in a round as a Product Expert, which depends on the Consumer's choice of *Purchase* or *Pass*.

Consumer's Decision on Purchasing Ranking Report	
Purchase	*Pass*
350	0

Table B.1: Product Expert's Potential Reward in ECU

Consumer's Reward. Your reward in a round consists of three parts:

(a) You will receive the value of the product you choose:

(i) 100 ECU if you have chosen Product A; and

(ii) either 0 or 80 ECU if you have chosen Product B, depending on the value randomly selected by the computer according to the 40%–60% chance.

(b) You will receive an extra reward of 250 ECU if you have chosen the first ranked product (regardless of whether you purchase the ranking report or not).

(c) You will pay a cost of 70 ECU if you choose to purchase the ranking report.

Table B.2 on the next page summarizes your potential reward in a round as a Consumer, which depends on

(a) your decision on purchasing the ranking report (*Purchase* or *Pass*);

(b) your product choice (*A* or *B*);

(c) the Product Expert's ranking method (*Method 1* or *Method 2*);

(d) chance (the randomly selected value of Product B); and

(e) the interaction between (c) and (d) which determines the generated report (*"Product A is Ranked 1st"* or *"Product B is Ranked 1st"*).

Product Expert's Choice of Ranking Method and the Generated Report

Your Decisions	*Method 1*	*Method 2*	
	"Product A is Ranked 1st" (always generated and seen by you)	*"Product A is Ranked 1st" (generated with 40% chance and seen by you)*	*"Product B is Ranked 1st" (generated with 60% chance and seen by you)*
Purchase & A	280	280	30
Purchase & B	−70 with 40% chance 10 with 60% chance	−70	260

Your Decisions	*Method 1*	*Method 2*	
	"Product A is Ranked 1st" (always generated and NOT seen by you)	*"Product A is Ranked 1st" (generated with 40% chance and NOT seen by you)*	*"Product B is Ranked 1st" (generated with 60% chance and NOT seen by you)*
Pass & A	350	350	100
Pass & B	0 with 40% chance 80 with 60% chance	0	330

Table B.2: Consumer's Potential Reward in ECU

The following three examples illustrate how to read and use the table.

(a) Scenario: (i) The Product Expert chooses *Method 2*; (ii) Report *"Product B is Ranked 1st"* is generated (i.e., 80 is selected for Product B's value); (iii) You choose to purchase the report (*Purchase*); and (iv) You choose *B* after seeing *"Product B is Ranked 1st."*

You will then receive: **80 (Product B's selected value) + 250 (Product B ranked 1st) – 70 (cost of report) = 260**.

(b) Scenario: (i) The Product Expert chooses *Method 2*; (ii) Report *"Product A is Ranked 1st"* is generated (i.e., 0 is selected for Product B's value); (iii) You choose not to purchase the report (*Pass*); and (iv) You choose *B* without viewing any report.

You will then receive: **0 (Product B's selected value) + 0 (Product B not ranked 1st) – 0 (not paying for report) = 0**.

(c) Scenario: (i) The Product Expert chooses *Method 1*; (ii) Report *"Product A is Ranked 1st"* is generated (always generated under *Method 1* irrespective of Product B's selected value); (iii) You choose not to purchase the report (*Pass*); and (iv) You choose *B* without viewing any report.

Case 1: If 0 is selected for Product B's value (40% chance), you will receive:
0 (Product B's selected value) + 0 (Product B not ranked 1st) – 0 (not paying for report) = 0.

Case 2: If 80 is selected for Product B's value (60% chance), you will receive:
80 (Product B's selected value) + 0 (Product B not ranked 1st) – 0 (not paying for report) = 80.

Information Feedback

At the end of each round, you will be provided with a summary of what happened in the round, including the Product Expert's ranking method, the selected value of Product B, the generated ranking report, the Consumer's decision on purchasing the ranking report, the Consumer's product choice, and your reward in the current round. A history of all previous rounds will also be provided.

Your Cash Payment

The experimenter randomly selects 3 rounds out of the 40 rounds to calculate your cash payment. So it is in your best interest to take each round equally seriously. The average of the ECU you earn in the 3 selected rounds will be converted into U.S. dollar at an exchange rate of 20 ECU for 1 USD. You will also separately receive a "show-up fee" of 5 USD.

Quiz and Practice

To ensure your understanding of the instructions, we will provide you with a quiz below. After the quiz, you will participate in a practice round. The practice round is part of the instructions and is not relevant to your earnings. Its objective is to get you familiar with the computer interface and the flow of the decisions in each round.

Once the practice round is over, the computer will ask you to "Click 'Next' to start the official rounds."

Quiz

To ensure your understanding of the instructions, we ask that you complete a short quiz before we move on to the experiment. This quiz is only intended to check your understanding of the written instructions. It will not affect your earnings. We will discuss the answers after you work on the quiz.

If you have any question, please raise your hand now. We will answer your question individually. If there is no question, we will proceed to the quiz.

1. True or False: The Product Expert makes the ranking-method decision after seeing which value the computer has selected for Product B. Circle one: True / False
2. True or False: The Product Expert has direct and immediate control over which ranking report, "Product A is Ranked 1st" or "Product B is Ranked 1st," is generated. Circle one: True / False
3. True or False: The generated ranking will not affect the Consumer's reward if he/she does not purchase the ranking report. Circle one: True / False
4. True or False: It costs the Consumer to view the generated ranking report.
   Circle one: True / False
5. True or False: The Consumer will see the ranking method chosen by the Product Expert if he/she purchases the ranking report. Circle one: True / False
6. True or False: The computer will directly reveal the selected value of Product B to the Consumer if he/she purchases the ranking report. Circle one: True / False