Shuffling as a Sales Tactic: An Experimental Study of Selling Product Rankings*

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

March 7, 2024

Abstract

We investigate the strategic interaction between a product expert and a consumer. The expert publicly commits to a ranking methodology to rank two products with uncertain relative merits; the consumer decides whether to acquire the 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. Strategic shuffling, in which the expert induces demand for his report by manipulating the uncertainty in product rankings, emerges as an equilibrium phenomenon. When the consumer highly values the top-ranked product, the expert-optimal equilibrium, which features shuffling, diverges from the consumer-optimal equilibrium. Laboratory evidence supports the predictions of the expert-optimal equilibrium. With limited field data due to proprietary ranking methodologies, our study provides useful alternative evidence on how ranking publishers may adopt methodologies that are not in consumers' best interests.

Keywords: Product Ranking; University Ranking; Product Guidance; Ranking Uncertainty; Laboratory Experiment

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

*We are grateful to Weijia (Daisy) Dai, David Goldbaum, Wenhao Wu, Oliver Yao, the department and associate editors, and two anonymous referees for their 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 National Science Foundation of China (Key Project No.: 72033006) and the Office of the Vice President and Associate Provost for Research and Graduate Studies at Lehigh University (Faculty Research Grant Account No.: 606975).

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

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

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

1 Introduction

Consumers often seek expert advice prior to making purchases, and the advice often takes the form of product rankings. Students and parents consulting university rankings (e.g., *U.S. News & World Report Best Colleges Ranking*), car buyers viewing auto rankings (e.g., *Kelley Blue Book Best Cars*), and home cooks accessing kitchen product rankings (e.g., *Cook's Illustrated*) are but few of the familiar examples. In each of these examples, a ranking publisher collects information about product attributes, chooses a ranking methodology that maps the attributes into a ranking, and offers the resulting ranking reports to consumers.

Product rankings provide informational guidance about product attributes. They may also influence consumer choices for reasons unrelated to product information.[1] Analogous to how an advertised product may be complementary to the advertisement itself due to imagine concerns (Becker and Murphy, 1993), a highly ranked product may confer sought-after social prestige. Product rankings thereby transform a product with n intrinsic attributes into one with $n+1$ attributes with an extra ranking attribute. This suggests that ranking publishers may influence the values of ranked products in a way orthogonal to how much the products are worth by themselves, reminiscent of a fashion magazine exercising the magic wand to dictate an otherwise unremarkable outfit as the season's stylish standard.

In enhancing the readership of its publication, a ranking publisher may put profits in front of consumer interests. Yet consumers' willingness to spend on accessing a ranking publication may be partly derived from the opportunity to learn about the contrived ranking attribute. Introspection would tell that this offers a perfect recipe for a principle-agent problem: Profit-driven ranking publishers may leverage their controls over ranking attributes to generate sales and traffic, even when doing so sacrifices the informational function of the rankings valuable to consumers. More than just an introspection, the problem has indeed been noted by observers of the industry; on university rankings, e.g., journalist Tierney (2013) wrote in *The Atlantic*:

> *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 profit motives as suggested, would have ramifications not only for consumers but also for other stakeholders such as university

[1] Pope (2009) and Luca and Smith (2013) find evidence in 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).

managements in the case of college rankings. Observations like the one above are, however, anecdotal. There is an inherent difficulty in obtaining direct empirical evidence given that ranking methodologies are often proprietary and complicated—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, some form of evidence beyond casual observations is needed. In this paper, we provide experimental evidence that a product expert, who benefits from a consumer's acquisition of his ranking advice amid its influence on product values, may adopt a ranking methodology that does not serve the best interests of the consumer.

We begin by analyzing a ranking-report game, which drives our experimental design and helps make precise the ideas expressed above. An expert (he), who cares only about whether a consumer (she) acquires his ranking report, publicly commits to a ranking method to generate the report. There are two products, and a ranking method is modeled as a probability distribution that the products are ranked first conditional on the values of their intrinsic attributes. The consumer, who is imperfectly informed about these intrinsic values, first decides whether to pay to view the ranking outcome and then chooses a product.

The consumer experiences the intrinsic value of the product of her choice and an additional ranking value if the product turns out to be ranked first. This ranking value accrues to the consumer even if she stumbles on the top-ranked product without the ranking report, and this represents a key feature of our environment: With this preference structure, the expert is in effect selling *product guidance* regarding intrinsic values as well as *resolution of ranking uncertainty* over which product carries the ranking value.

A crucial insight of our equilibrium analysis centers on a phenomenon that we term "shuffling as a sales tactic." Leveraging the effect of resolving ranking uncertainty provided by his report, the expert engages in *strategic shuffling*—endogenously manipulating the uncertainty by choosing a method that sometimes ranks the intrinsically less valuable product first—to induce demand, even when doing so misguides the consumer in respect of the intrinsic values. The consumer is willing to pay for the report, enduring the poor guidance, because the shuffling makes it even worse to do without it.

This apparently paradoxical situation is most palpable to the consumer when the ranking value is relatively high, in which the expert-optimal equilibrium diverges from the consumer-optimal equilibrium. In the expert-optimal equilibrium, the consumer's willingness to pay for the report is maximal—it is as if the expert created a problem by shuffling and then peddled the solution. In the consumer-optimal equilibrium, the consumer's expected payoff from acquiring the report is instead maximal. When the ranking value is relatively low so that the resolution of ranking uncertainty is not that valuable, the two equilibria coincide.

We simplify the game for experimental implementation and conduct four treatments with variations in the ranking value and the cost of the report to the consumer. Two treatments belong to the case where the expert-optimal and the consumer-optimal equilibria coincide and the other two the case where they diverge.

Our experimental findings support the qualitative predictions of perfect Bayesian equilibria. For experts, the ranking methods under which consumers are predicted to acquire reports are on average chosen more often than those under which consumers do not acquire. For consumers, their report-acquisition decisions are overall in lockstep with experts' choices, and the products chosen with and without reports reflect sequential rationality.

A more demanding and interesting test of the theory lies in the competition between the expert-optimal and the consumer-optimal equilibria. Our data favor the former. In the treatments where the predictions of the two equilibria diverge, the ranking method that is expert-optimal and features shuffling is most frequently chosen by a considerable margin. A panel data analysis reveals a more refined picture: When a ranking method that is both consumer-optimal and expert-optimal becomes not expert-optimal in another treatment, the method is significantly less likely to be chosen; by contrast, an expert-optimal ranking method is significantly more likely to be chosen even when it is not consumer-optimal.

Our theoretical and empirical findings shed light on the sentiment about ranking publishers that motivates our study. An expert who deliberately generates a probabilistic ranking outcome to induce demand parallels a ranking publisher who frequently tweaks its rankings to draw consumers to follow its publications. The possible divergence between the expert-optimal and the consumer-optimal equilibria furnishes a theoretical basis to argue that the tweaking can be excessive from the vantage point of consumer welfare. On the empirical side, the prevalence of the expert-optimal equilibrium in a low-stake experimental setting lends credence to the real-world presence of the principal-agent problem: It stands to reason that outside the laboratory with substantially higher stakes, ranking publishers operating under comparable incentive structures might indeed be putting profits in front of consumer welfare.

Related Literature. Our study relates to two separate strands of literature: in terms of the subject matter, product rankings and more generally non-seller-provided product information; in terms of the theory and experiment, strategic information transmission.

In respect of product rankings, our game shares common features with the university-ranking model of Dearden et al. (2019). In their dynamic model, universities with finite numbers of attributes are ranked in each period by a ranking publisher. The publisher chooses the weight of each attribute and aggregates them into a ranking. As in our case, students derive prestige utilities from attending the top-ranked universities, and the publisher injects uncertainty into its ranking methodology. Our static model with a different type of methodology captures the same key result while providing a simpler environment for experimental inquiry.

A ranking publisher selecting the top-ranked products parallels a fashion magazine editor dictating a season's "it" outfits, which sometimes also appear to done randomly. Kuksov and Wang (2013) analyze a model where stylish consumers, who prefer to be identified, have exclusive access to the product recommendations of a coordinator interpreted as a fashion magazine. By following the random recommendations, these high-type stylish consumers separate themselves from the low types. The seeming randomness in product rankings in our case and fashion hits in theirs are commonly rationalized as outcomes of maximizing behavior.

Recommendations by product magazines or ranking publishers are not the only non-seller sources of product information. Online product reviews submitted by consumers, e.g., in physician ratings (Lu and Rui, 2018), have been shown to provide useful product information. Yet the ubiquity of fake reviews, either submitted by sellers themselves or competitors (Mayzlin et al., 2014), dilutes the value of information (Anderson and Simester, 2014), to the extent that some customer review 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 also be driven by profit motives to provide misguiding product information.

Our ranking-report game shares the key features of games of strategic information transmission: A sender with private information influences the action of a receiver via a message. Our setting, however, differs from the three canonical environments of this genre with payoff-dependent (costly) messages (Spence, 1973), payoff-independent (cheap-talk) messages (Crawford and Sobel, 1982), and state-dependent (verifiable) messages (Grossman, 1981; Milgrom, 1981). The message in our game—the ranking report—directly influences payoff but it is the payoff of the receiver (consumer). Furthermore, our expert provides comparative information on ordinal rankings and takes no direct interest in the action (product choice) of the consumer.[2]

Another critical distinction of our game is that the information is not transmitted as a direct execution of strategy; rather, it is done under a committed signaling rule (the ranking method). This aspect of our environment is closer to the recent literature on Bayesian persuasion pioneered by Kamenica and Gentzkow (2011). Unlike our game, however, there is no element of selling information in Bayesian persuasion.[3]

Given the theoretical connection, our experiment is naturally related to the experimental

[2]There are extensions of the canonical cheap-talk games that share some of these features. Chakraborty and Harbaugh (2007) study the transmission of comparative information that takes the form of rankings of multidimensional issues. Ottaviani and Sørensen (2006) consider a sender who takes interest only in being perceived as informed in front of a receiver, not in any explicit action that the receiver may take.

[3]It is worth comparing the roles of 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., 2023) have shown that random message transmissions could improve cheap-talk communication, resulting in Pareto improvements. By contrast, randomness in our case benefits one party but harms the other. Additionally, in those games an exogenous communication protocol generates the randomness, while our random reports are partly endogenous.

strand of the literature on information transmission. Recent work includes Schmidt and Buell (2017) and Fudenberg and Vespa (2019) for costly signaling, Gneezy (2005), de Groot Ruiz et al. (2015), and Lai and Lim (2018) for cheap talk, Hagenbach and Perez-Richet (2018) and Jin et al. (2021) for verifiable disclosures, and Fréchette et al. (2022), Wu and Ye (2023), and Au et al. (2023) for Bayesian persuasion. Choosing posteriors lies at the heart of Bayesian persuasion. An experimental focus has been on how to assist subjects to accurately update beliefs. Our experimental design also contributes to this endeavor.

Section 2 analyzes our experimental ranking-report game. Section 3 describes our experimental design and hypotheses. We report our findings in Section 4. Section 5 concludes.

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 products and derives benefit if the consumer acquires the resulting ranking *report*. The imperfectly informed consumer makes two decisions, whether to acquire the report and which product to choose.

Consumer Utility. The consumer derives utility from the intrinsic attributes and the ranking attribute of the chosen product. Intrinsic attributes, which include exogenously determined quality, price, etc., give rise to *intrinsic values*. 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 and equals either 0 or $\bar{v}_B > 0$. The common prior is that $v_B = \bar{v}_B$ with probability $0 < p < 1$.

Ranking attributes are determined through the expert's endogenous choice of ranking method. A product yields to the consumer $r > 0$ if it is ranked first and 0 otherwise, independent of its intrinsic value. We call r the *ranking value* of the top-ranked product.

We impose two parameter restrictions, where the first simplifies the cases for expositional brevity and the second ensures that the ranking reports influence product choices:

Assumption 1 (Parameters)**.** *The parameters satisfy (a)* $\bar{v}_B > \bar{v}_A \neq r$ *and (b)* $p\bar{v}_B + r - \bar{v}_A > 0$.

Ranking Methods. The expert chooses and commits to a ranking method, which is a mapping $\beta : \{0, \bar{v}_B\} \to [0, 1]$ specifying for each possible intrinsic value of Product B a probability that Report B is issued (Report K ranks Product K first, $K = A, B$). To design a parsimonious experimental game, we restrict attention to the class of ranking methods where $\beta(\bar{v}_B) = 1$. The expert's choice thereby reduces to choosing $\beta(0) = \beta_0 \in [0, 1]$.[4]

[4]The set of equilibrium ranking methods of this simpler game is a subset of that of the general game where

The consumer observes β_0 but can access the ranking report if and only if she pays an exogenously given fee $f > 0$.[5] With or without viewing the report, the consumer then chooses between Products A and B (product prices are part of the intrinsic values).

Our experiment explores the interplay of two properties of ranking methods. The first concerns *product guidance*, which is provided if the more intrinsically valuable product is ranked first.[6] Product guidance is always provided when $v_B = \bar{v}_B$. When $v_B = 0$, Report A guides the consumer while Report B misguides, and $1 - \beta_0$ measures a ranking method's likelihood of offering guidance. This gives rise to the following property:

Fact 1. *Product guidance worsens as β_0 increases.*

The second property concerns *ranking uncertainty*. While a report reveals the ranking, a ranking method may induce uncertainty about it. Without the report, the consumer is typically unsure of which product carries the ranking value. The prior p represents the "natural level" of this uncertainty. By engaging in *strategic shuffling* with $\beta_0 > 0$, the expert can endogenously manipulate the uncertainty to deviate from its natural level. The following characterization of β_0, p, and the ranking uncertainty informs our choice of the relevant experimental parameter:

Fact 2. *Ranking method $\beta_0 = 0$ is uncertainty-neutral (not altering the natural uncertainty), whereas $\beta_0 = 1$ is uncertainty-eliminating. For $p < \frac{1}{2}$, $\beta_0 \in \left(0, \frac{1-2p}{1-p}\right)$ adds on to and $\beta_0 \in \left(\frac{1-2p}{1-p}, 1\right)$ suppresses the natural uncertainty. For $p \geq \frac{1}{2}$, any $\beta_0 > 0$ suppresses the natural uncertainty.*[7]

The ability to control the ranking uncertainty acts as the expert's strategic instrument to induce demand for his uncertainty-resolving ranking report.

Strategies and Payoffs. A pure strategy of the expert is a choice $\beta_0 \in [0, 1]$. A pure strategy of the consumer comprises: (a) a report-acquisition decision rule, $s : [0, 1] \to \{0, 1\}$, specifying for each β_0 whether to acquire the report ($s = 1$) or not ($s = 0$), and (b) a product choice rule, $a : \{A, B, \varnothing\} \to \{0, 1\}$, specifying for given (β_0, s) whether to choose Product A ($a = 1$) or B ($a = 0$) after viewing report A, B, or none ($\varnothing$). We call $a(\varnothing)$ the *default product choice*.[8]

$\beta(\bar{v}_B) = 1$ is relaxed. Furthermore, the selected equilibria (Section 2.3) coincide in both games. The analysis of the general game is available upon request.

[5] The fee can be interpreted as a subscription fee or more generally any cost that the consumer incurs to access the report (e.g., search cost or cost of attention), not necessarily a payment to the expert.

[6] Product guidance so defined is an information about the ordinal rankings of intrinsic values. To cover cardinal values, one may extend our model with a richer space of values and finer ranking categories such as "Product X is hands down better than Product Y" and "Product X is slightly better than Product Y."

[7] When the probabilities of the two products being ranked top become more uniform, ranking uncertainty rises as measured by entropy (Shannon, 1948). Shuffling is deemed to exist when the ranking uncertainty deviates from its natural level, including the degenerate $\beta_0 = 1$ which eliminates the natural uncertainty.

[8] A complete contingency plan specifies the consumer's product choice in every subgame associated with a β_0 and s. For notational brevity, we omit β_0 as an argument of a, while s is captured by A, B, and $\varnothing$.

We focus on pure strategies, which substantially simplifies the exposition with minimal loss of generality. For the expert, a mixed strategy is analogous to a compound lottery reducible to a simple lottery. For the consumer, we assume a tie-breaking rule, which spares us from wordy statements of randomization for the knife-edge cases and from dwelling on the non-critical details of whether the equilibrium ranges of β_0 include the endpoints or not:

Assumption 2 (Tie Breaking)**.** *The consumer's indifference is resolved in favor of, for report acquisition, acquiring the ranking report, and, for product choice, choosing Product A.*

The expert's payoff equals the revenue $\pi > 0$ derived from the consumer's acquisition of his report; his choice of ranking method and the consumer's product choice do not directly affect his payoff. While the expert's revenue is not necessarily the same as the consumer's payment, for expositional convenience we refer to the expert selling the report to the consumer.[9]

The consumer's payoff equals her utility from the chosen product (the intrinsic value plus any ranking value) less any report fee f. Given ranking outcome $\mathbb{I}_A \in \{0,1\}$, where $\mathbb{I}_A = 1$ indicates that Product A is ranked first, a consumer who makes report-acquisition decision $s \in \{0,1\}$ and product choice $a = \{0,1\}$ with resulting intrinsic value $v(a) \in \{\bar{v}_A, v_B\}$ receives:

$$u(s,a,v(a),\mathbb{I}_A) = \begin{cases} v_B + r(1-\mathbb{I}_A) - sf & \text{if } a = 0, \\ \bar{v}_A + r\mathbb{I}_A - sf & \text{if } a = 1, \end{cases}$$

where v_B is drawn to be either 0 or $\bar{v}_B$ according to the prior p.

The specification that the consumer earns r for choosing the top-ranked product whether or not she acquires the report is a key feature of our game. It allows the expert's choice of β_0 to influence how much the resolution of ranking uncertainty is worth to her.[10]

2.2 Equilibrium Characterization

Under Assumption 1(b), the consumer always chooses the top-ranked product after viewing the ranking report, i.e., $a(A) = 1$ and $a(B) = 0$. While this is necessary for the consumer to be willing to pay for the report in equilibrium, it is not sufficient. The alternative to acquiring the report—the default product—also plays a role. The following lemma characterizes the optimal default product choice in terms of β_0 and a threshold $\beta_{AB} = \frac{\bar{v}_A - p\bar{v}_B + (1-2p)r}{2(1-p)r}$:

[9] In practice, ranking publishers typically have more lucrative income sources that are incidental to their publications. In a two-sided market, even if they offer their product rankings to consumers for free, they derive revenues from advertising or consulting services provided to product sellers.

[10] This also reflects the effects of rankings in reality. A student attending the top-ranked university may enjoy the prestige regardless of whether the attending decision is made with consulting ranking publications or not. The ranking value may also take the form of future economic value: A highly ranked car model is likely to command a higher resale value regardless of whether the owner learns about its ranking or not.

Lemma 1. *Product A is the optimal default product, i.e., $a(\varnothing) = 1$, if and only if $\beta_0 \leq \beta_{AB}$.*

The consumer's expected utility evaluated after she observes β_0 but before viewing any report depends on and, in turn, shapes her default product choice and report-acquisition decision. If she eventually does not acquire the report, then this expected utility equals that from a default product. By Lemma 1, her expected utility from not acquiring the report and optimally choosing a default product is thus:

$$V^{\varnothing}(\beta_0) = \begin{cases} V_A^{\varnothing}(\beta_0) & \text{if } \beta_0 \leq \beta_{AB}, \\ V_B^{\varnothing}(\beta_0) & \text{if } \beta_0 > \beta_{AB}, \end{cases} \tag{1}$$

where $V_A^{\varnothing}(\beta_0) = \bar{v}_A + (1-p)(1-\beta_0)r$ and $V_B^{\varnothing}(\beta_0) = p\bar{v}_B + [p + (1-p)\beta_0]r$ are the respective expected utilities from default Products A and B, with β_{AB} satisfying $V_A^{\varnothing}(\beta_{AB}) = V_B^{\varnothing}(\beta_{AB})$.

If the consumer eventually acquires the report, then her expected utility before viewing the report is:

$$V(\beta_0) = p\bar{v}_B + (1-p)(1-\beta_0)\bar{v}_A + r. \tag{2}$$

The consumer's report-acquisition decision depends on whether the expected utility gain $G(\beta_0) = V(\beta_0) - V^{\varnothing}(\beta_0)$—her *willingness to pay* for the report—measures up to the report fee f. The expert can control this willingness to pay through his choice of a ranking method. The following proposition characterizes the equilibrium choices of β_0:[11]

Proposition 1. *In any pure-strategy perfect Bayesian equilibrium of the game, the expert sells the ranking report by choosing a β_0 such that $G(\beta_0) \geq f$; in the event that $G(\beta_0) < f$ for all $\beta_0 \in [0,1]$, the expert chooses any β_0 without selling the report.*

The equilibria in which the consumer acquires the ranking report form the focus of our experimental inquiry. In these *acquisition equilibria*, the expert chooses a β_0 that renders the consumer's willingness to pay greater than the moderate report fee.

Four cases of acquisition equilibria, which underscore the tension between product guidance and ranking uncertainty, will inform our treatment design. Let $f_1 = p(\bar{v}_B + r - \bar{v}_A)$ and $f_2 = \frac{(p\bar{v}_B + r - \bar{v}_A)(\bar{v}_A + r)}{2r}$. Table 1 lists the equilibrium ranges of β_0 in the four cases in terms of $\beta_A = \frac{f - f_1}{(1-p)(r - \bar{v}_A)}$ and $\beta_B = 1 - \frac{f}{(1-p)(\bar{v}_A + r)}$, where β_A satisfies $V(\beta_A) - f = V_A^{\varnothing}(\beta_A)$ and β_B satisfies $V(\beta_B) - f = V_B^{\varnothing}(\beta_B)$. The cases are categorized by the different levels of report fee relative to

[11]The expert's choice of β_0 initiates a subgame, in which the consumer updates beliefs in choosing a product. To impose sequential rationality on the consumer on and off equilibrium paths with beliefs derived from Bayes' rule whenever possible, we use perfect Bayesian equilibrium as the solution concept.

Table 1: Equilibrium Ranking Methods

	$f \in (0, f_2]$	$f \in (f_2, f_1]$
$r < \bar{v}_A$	**Case 1** $\beta_0 \in [0, \beta_B]$	**Case 2** $\beta_0 \in [0, \beta_A]$

	$f \in (0, f_1]$	$f \in (f_1, f_2]$
$r > \bar{v}_A$	**Case 3** $\beta_0 \in [0, \beta_B]$	**Case 4** $\beta_0 \in [\beta_A, \beta_B]$

f_1 and f_2. The key dichotomy concerns whether the effect of product guidance is more valuable than the effect of resolving ranking uncertainty ($r < \bar{v}_A$) or vice versa ($r > \bar{v}_A$).[12]

The expected utility gain $G(\beta_0) = V(\beta_0) - V^{\varnothing}(\beta_0)$ embodies both effects. The effect of product guidance diminishes as β_0 increases (Fact 1), which manifests as $V(\beta_0)$ in (2) being decreasing in β_0. To disentangle the effect of resolving ranking uncertainty, we may hypothetically hold the guidance effect constant by setting $\bar{v}_A = \bar{v}_B = 0$ so that $V(\beta_0)$ is constant at r. This reveals that, with $V^{\varnothing}(\beta_0)$ in (1) being a V-shaped upper envelope, the value of resolving uncertainty, $r - V^{\varnothing}(\beta_0)$, increases and then decreases in β_0, attaining the maximum at β_{AB}.[13]

With these observations, the cases of equilibria can be exemplified with Figure 1. Panel (a) illustrates Case 1 where, with $r < \bar{v}_A$, product guidance is more valuable than resolving ranking uncertainty. The curve $V(\beta_0) - f$ is steep relative to $V^{\varnothing}(\beta_0)$, with $G(\beta_0) - f$ strictly decreasing in β_0. Product guidance is the paramount factor when r is low, and the consumer acquires the report for low enough β_0 with sufficient product guidance. The fee f pins down the intersection of $V(\beta_0) - f$ and $V^{\varnothing}(\beta_0)$ beyond which the report becomes not worthwhile. The intersection occurs at β_B in Case 1 with a relatively low f. Had f been higher as in Case 2, $V(\beta_0) - f$ would have shifted down and intersected with $V^{\varnothing}(\beta_0)$ at a β_A leftward of β_{AB}.

Panel (b) illustrates Case 4, where, with $r > \bar{v}_A$, resolving ranking uncertainty is more valuable than product guidance. The curve $V(\beta_0) - f$ is now flat relative to $V^{\varnothing}(\beta_0)$, and $G(\beta_0) - f$ is inverted V-shaped peaking at β_{AB}. With resolving ranking uncertainty being the paramount factor when r is high, the consumer acquires the report for β_0 in the neighborhood of where this resolution is most valuable. In the illustrated Case 4 with a relatively high f, this neighborhood is enclosed by β_A and β_B. Had f been lower as in Case 3, the neighborhood would have extended all the way on the left to include 0.

[12] Recall that $r = \bar{v}_A$ is ruled out by Assumption 1(a). For $r < \bar{v}_A$, $f_2 \le f_1$, while for $r > \bar{v}_A$, $f_1 \le f_2$. Note also that in all four cases both Products A and B have the potential to be the optimal default ($0 \le \beta_{AB} < 1$).

[13] We thank an anonymous referee for pointing out these observations, which make our exposition considerably more succinct.

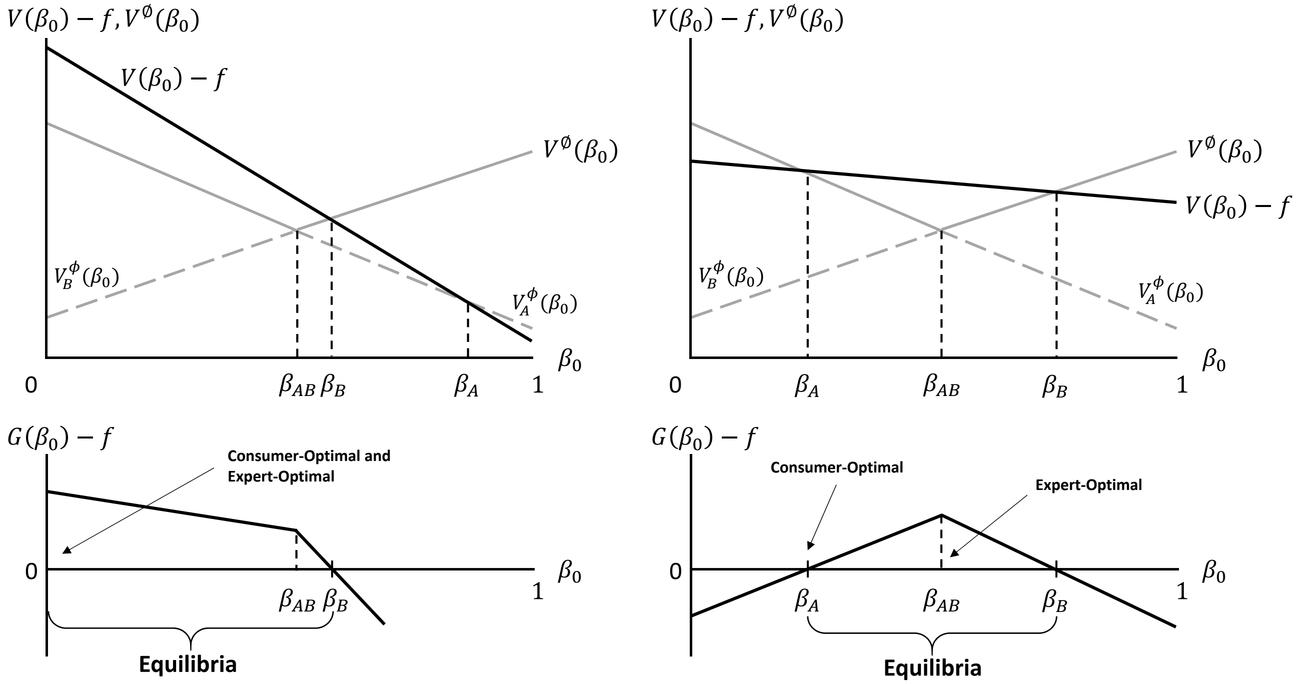

(a) $r < \bar{v}_A$ and $f \in (0, f_2]$ (b) $r > \bar{v}_A$ and $f \in (f_1, f_2]$

Figure 1: Consumer's Expected Payoffs and Equilibrium Ranking Methods

2.3 Equilibrium Selection

The above lays down the comparative statics with respect to the ranking value r and report fee f. There is nonetheless a wide range of ranking methods that survive the restrictions. We further restrict behavior by selecting the *consumer-optimal* and the *expert-optimal* equilibria.

The expert's choice of ranking method can be seen as an effort to make the report alluring to the consumer, and it is intuitive to select based on the most alluring report. The consumer-optimal equilibrium corresponds to one interpretation of "most alluring": one that renders the *acquiring* consumer the highest expected payoff $V(\beta_0) - f$. The expert-optimal equilibrium corresponds to another interpretation: one that imposes the highest expected deviating loss $G(\beta_0) - f$ on the *non-acquiring* consumer, which coincides with the highest willingness to pay.

Table 2 lists the selected equilibrium β_0 among the set of acquisition equilibria in each of the four cases. In addition to the intuitive appeal, the two selected equilibria can also be singled out by formal criteria. Efficiency would select the consumer-optimal equilibrium.[14] A perturbation-based refinement in the spirit of Myerson's (1978) proper equilibrium would select the expert-optimal equilibrium. We fully develop this refinement in Appendix A.2. It provides

[14]Since the expert receives π in any acquisition equilibrium, the equilibria can be Pareto ranked by the consumer's expected payoff. Note that any "ϵ-altruism" of the expert toward the consumer would also select the consumer-optimal equilibrium. In an experimental study of coordination games, Chen and Chen (2011) use group-identity based altruistic preferences to select the more efficient equilibria.

Table 2: Selected Equilibrium Ranking Methods

	$f \in (0, f_2]$	$f \in (f_2, f_1]$
$r < \bar{v}_A$	**Case 1** Consumer-Optimal $\beta_0 = 0$ Expert-Optimal $\beta_0 = 0$	**Case 2** Consumer-Optimal $\beta_0 = 0$ Expert-Optimal $\beta_0 = 0$

	$f \in (0, f_1]$	$f \in (f_1, f_2]$
$r > \bar{v}_A$	**Case 3** Consumer-Optimal $\beta_0 = 0$ Expert-Optimal $\beta_0 = \beta_{AB}$	**Case 4** Consumer-Optimal $\beta_0 = \beta_A$ Expert-Optimal $\beta_0 = \beta_{AB}$

a more precise sense for the equilibrium to be expert optimal: When the consumer trembles, the report is most likely to be acquired when the consumer's willingness to pay is the highest.

We conclude our theoretical analysis by examining the coincidence and discrepancy of the two selected equilibria through the lens of product guidance and ranking uncertainty. A key observation is that manipulation of the ranking uncertainty beyond what is necessary to induce report acquisition harms the acquiring consumer by worsening the product guidance, with no marginal benefit for the expert. Consequently, the consumer-optimal equilibrium admits the minimum β_0 subject to the report being acquired, i.e., $G(\beta_0) - f \geq 0$. Panels (a) and (b) of Figure 1 illustrate that this constrained minimum occurs at, respectively, $\beta_0 = 0$ in Case 1 (the same for Cases 2 and 3) and $\beta_0 = \beta_A$ in Case 4.

Turning to the expert-optimal equilibrium, another key observation is that while the uncertainty manipulation harms the acquiring consumer, it could enhance willingness to pay, which is the payoff difference between the acquiring and non-acquiring consumer shaped by both product guidance and ranking uncertainty. When the consumer values guidance more than uncertainty resolution ($r < \bar{v}_A$), however, the same consideration of maximal product guidance that drives the consumer-optimal equilibrium dominates; the two selected equilibria coincide and commonly admit $\beta_0 = 0$ (Cases 1 and 2). When instead the consumer values uncertainty resolution more ($r > \bar{v}_A$), the expert-optimal equilibrium admits $\beta_0 = \beta_{AB}$ that involves more manipulation and less guidance than the $\beta_0 = 0$ (Case 3) or $\beta_0 = \beta_A$ (Case 4) of the consumer-optimal equilibrium; the acquiring consumer is harmed, but the non-acquiring incarnation is harmed even more, making the solution peddled by the expert—the ranking report—most worthwhile to acquire in the expert-optimal equilibrium.[15]

[15]With the report fee fixed, the expert does not benefit directly from a higher willingness to pay except in the perturbed game of our formal refinement (Appendix A.2). In an alternative model in which f is the expert's endogenous choice, the fee will be set at the maximum willingness to pay, and the expert-optimal equilibrium

3 Experimental Implementation

3.1 Treatment Parameters

The four cases of acquisition equilibria yield *four treatments*. We adopt the same parameters $p = 0.2$, $\bar{v}_A = 100$, $\bar{v}_B = 250$, and $\pi = 300$ for all treatments and induce variations in r and f as follows: $(r, f) \in \{(55, 5), (55, 30), (250, 5), (250, 110)\}$. We choose these values to induce salient incentives with large payoff differentials. The use of 0.2 for the prior is informed by Fact 2, where a value less than $\frac{1}{2}$ renders a rich environment with the natural ranking uncertainty manipulable in either direction.

Table 3: Treatment Parameters and Theoretical Predictions

	Low Report Fee ($f = 5$)	Medium Report Fee ($f = 30$)
Low Ranking Value ($r = 55$)	***LL*** **Case 1** Acquisition: $\beta_0 \in \left[0, \frac{119}{124}\right]$ Consumer-Optimal: $\beta_0 = 0$ Expert-Optimal: $\beta_0 = 0$	***LM*** **Case 2** Acquisition: $\beta_0 \in \left[0, \frac{11}{36}\right]$ Consumer-Optimal: $\beta_0 = 0$ Expert-Optimal: $\beta_0 = 0$

	Low Report Fee ($f = 5$)	High Report Fee ($f = 110$)
High Ranking Value ($r = 250$)	***HL*** **Case 3** Acquisition: $\beta_0 \in \left[0, \frac{55}{56}\right]$ Consumer-Optimal: $\beta_0 = 0$ Expert-Optimal: $\beta_0 = \frac{1}{2}$	***HH*** **Case 4** Acquisition: $\beta_0 \in \left[\frac{1}{4}, \frac{17}{28}\right]$ Consumer-Optimal: $\beta_0 = \frac{1}{4}$ Expert-Optimal: $\beta_0 = \frac{1}{2}$

Note: All treatments share the parameter values $p = 0.2$, $\bar{v}_A = 100$, $\bar{v}_B = 250$, and $\pi = 300$.

We label the four treatments by *LL*, *LM*, *HL*, and *HH*, where the first letter refers to Low (55) or High (250) for the ranking value and the second letter refers to Low (5), Medium (30), or High (110) for the report fee.[16] Table 3 lists the treatments and the corresponding predictions of the acquisition, consumer-optimal, and expert-optimal equilibria.

will be the unique acquisition equilibrium.

[16] We strictly follow the acquisition equilibria in designing our treatments. The *LH* combination $(r, f) = (55, 110)$ is not included because it admits only non-acquisition equilibria, and the *HM* combination $(r, f) = (250, 30)$ admits the same class of equilibria as $(r, f) = (250, 5)$ but with smaller payoff differentials.

3.2 Design and Procedures

Our experiment was conducted using oTree (Chen et al., 2016) at the Experimental Economics Laboratory of Southwestern University of Finance and Economics. A total of 316 undergraduate subjects with no prior experience in the experiment participated. Upon arrival, each subject received a copy of a summary of the experimental instructions. The summary was read aloud by the experimenter, and subjects were given time to go through the more detailed on-screen instructions before they completed a comprehension quiz and a practice round.[17]

Four sessions were conducted per treatment with 18 – 24 subjects per session. A *between-subject* design was used. Half of the subjects were randomly assigned to the role of an expert and the other half the role of a consumer. Roles remained fixed throughout a session. Experts and consumers formed groups of two to play *40 rounds* of the game under *random matchings*.

We balance a faithful implementation of the game with a design that is conducive to the comprehension of the problem. The latter is particularly important given that choosing a ranking method is tantamount to choosing the consumer's posteriors, and subjects are notoriously non-Bayesian (e.g., Camerer, 1995). We settle with two design choices: (a) for simplicity, we discretize the set of ranking methods into *five choices*, and (b) for user-friendliness, we present joint probabilities graphically to help subjects process conditional probabilities.

All treatment parameters other than the prior were induced as monetary incentives denominated in Experimental Currency Unit (ECU). The fixed value of Product A was 100 ECU. The uncertain value of Product B was either 0 ECU with 80% chance or 250 ECU with 20% chance, drawn by the computer in each round. Product A (B) was referred to subjects as the better product if the value of Product B was 0 (250) ECU.

The expert made one decision in each round, choosing one of five ranking methods: 0%, 25%, 50%, 75%, and 100%, referred to as Methods 1, 2, 3, 4, and 5 respectively. All the methods always ranked Product B first when it was the better product. When Product A was the better product, they ranked Product B first with the corresponding percentages. These percentages cover the consumer-optimal and the expert-optimal equilibrium β_0 in all four treatments.

To help subjects visualize the conditional probabilities, the joint probability of each product being the better product and ranked first under each ranking method was depicted using the bar chart in Figure 2(a). The chart was displayed on the expert's decision screen.

After the expert selected a ranking method, the consumer made the first of two decisions. It was emphasized that the expert did not know which product was better when choosing a

[17] The experiment was conducted in Chinese. We first composed the experimental instructions in English and then translated them into Chinese. The sample English instructions for treatment *HH*, both the English and the Chinese versions, can be found in the online appendices.

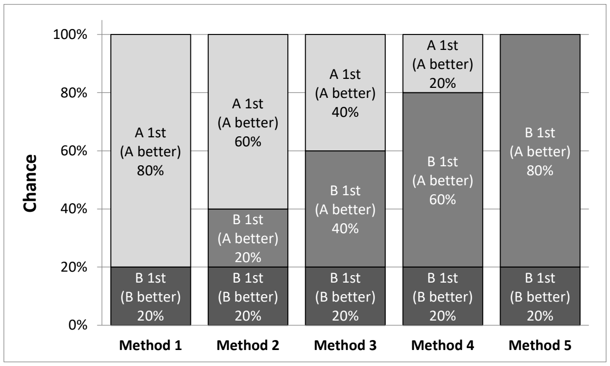

(a) Five Methods

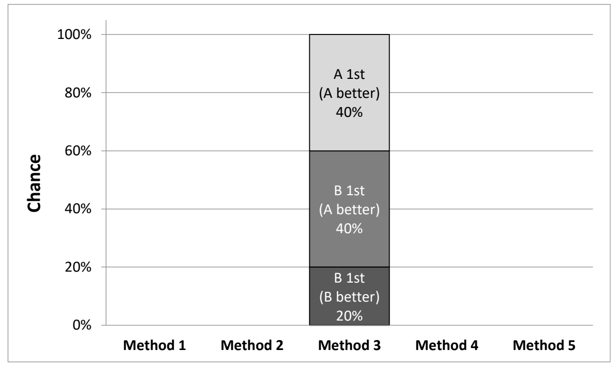

(b) A Selected Method

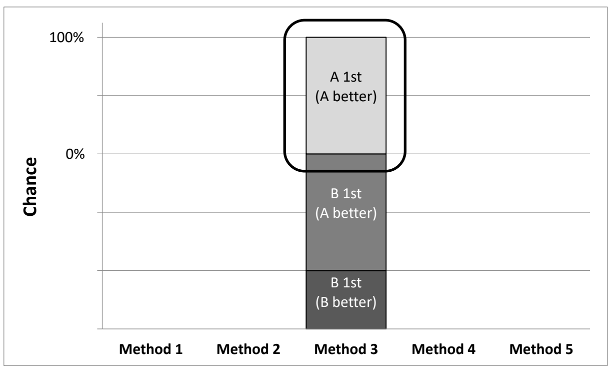

(c) Product A Ranked First

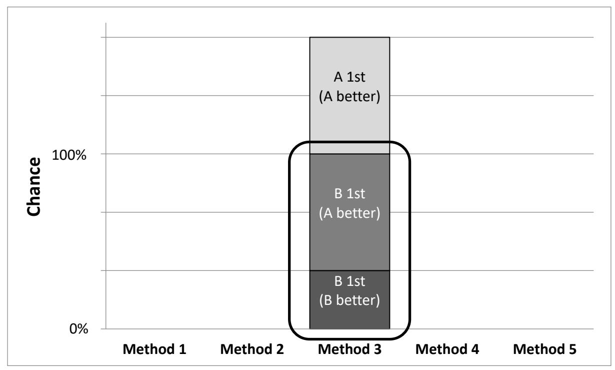

(d) Product B Ranked First

Figure 2: Ranking Methods and Reports Presented to Subjects

method. The selected method was revealed to the consumer in a similar bar chart. Figure 2(b) shows an example where Method 3 was selected. The consumer then decided whether to pay 110 ECU (treatment *HH* for an illustration) to learn which product was ranked first.

If the consumer decided not to pay for the ranking report, then the bar char in Figure 2(b) would remain on the screen. If the consumer instead decided to pay, then, depending on the draw of the value of Product B and thus which product was ranked first, one of the two charts in panel (c) or (d) would be shown. With or without the report, the consumer then chose a product, which concluded the round.

The expert would earn 300 ECU if the consumer acquired the report (nothing otherwise). The consumer earned the value of the chosen product, where the first ranked product was worth an extra 250 ECU (treatment *HH*). The earning from the product would be deducted by the report fee if the consumer paid for the report.

An information feedback was provided for each round. It covered the expert's choice of ranking method, the realized value of Product B, the top-ranked product, the consumer's report-acquisition decision and product choice, and the subject's earning for the round.

We randomly selected three out of the 40 rounds for calculating subject payments. The

Table 4: Predicted Ranking Methods

Treatment	(1) Acquisition	(2) Consumer-Optimal	(3) Expert-Optimal
LL	Methods $1, 2, 3, 4$	Method 1	Method 1
LM	Methods $1, 2$	Method 1	Method 1
HL	Methods $1, 2, 3, 4$	Method 1	Method 3
HH	Methods $2, 3$	Method 2	Method 3

Note: Columns (1), (2), and (3) contain the ranking method(s) predicted respectively by the acquisition, consumer-optimal, and expert-optimal equilibria among Method 1 (0%), Method 2 (25%), Method 3 (50%), Method 4 (75%), and Method 5 (100%).

average ECU earned in the three selected rounds was converted into Chinese RMB at a fixed and known exchange rate of 4 ECU for 1 RMB. A show-up fee of 20 RMB was also paid. A session lasted about an hour, and subjects on average earned 62.02 RMB.[18]

3.3 Experimental Hypotheses

Table 4 adapts the theoretical predictions summarized in Table 3 to the five experimental ranking methods and forms the basis for our experimental hypotheses on expert behavior. The discretization of ranking methods reduces the continuum of subgames initiated by $\beta_0 \in [0, 1]$ down to five and entails no change in the logic of equilibria.

The acquisition, consumer-optimal, and expert-optimal equilibria as three equilibrium concepts give rise to separate sets of testable comparisons, some of which are mutually exclusive and present competing hypotheses. We conduct within- and between-treatment comparisons for each concept. As experimental data are often noisy, the comparisons are qualitative. The general statements of hypothesis that are portable for the three equilibrium concepts are:

(a) Within a treatment, the average relative frequency of equilibrium ranking methods is greater than that of the non-equilibrium methods.

(b) Between two treatments, the relative frequency of a ranking method is greater in the treatment in which it is equilibrium than in the treatment in which it is not.

The detailed comparisons depend on which equilibrium concept is under consideration. Take the predictions of acquisition equilibria in column (1) of Table 4 as an example. The within-treatment hypothesis for, say, HH is that the average relative frequency of Methods 2

[18]As a point of reference, the hourly minimum wage in Beijing was RMB 25.3 in 2021, which was the highest among all regions in China.

and 3 is greater than that of the average relative frequency of Methods 1, 4, and 5.[19] The between-treatment hypothesis for, say, *HH* and *HL* is that, under the treatment effect of a higher report fee, the average relative frequency of Methods 1 and 4 is lower in *HH*.

Columns (2) and (3) of Table 4 suggest that some of the hypotheses for the consumer-optimal and the expert-optimal equilibria are competing. For treatments with low ranking value (LL and LM), both concepts make the same within-treatment predictions: Method 1 is the modal method. For treatments with high ranking value (HL and HH), however, the two concepts predict differently: The expert-optimal equilibrium predicts Method 3 to be modal, while the consumer-optimal equilibrium predicts either Method 1 or Method 2. This contrast sets apart the equilibrium concepts; the high-ranking-value treatments are our principal focus to empirically evaluate and differentiate the two selected equilibria.

For consumers, the general hypothesis is again that equilibrium or optimal behavior is observed more often. In the theory, the expert never chooses in an acquisition equilibrium a ranking method where the consumer does not acquire the report; the consumer's rationality off the equilibrium path is never subject to test. In the experiment, we expect all methods to be chosen as part of noisy behavior, and that allows us to see if consumers behave as predicted by sequential rationality, even in cases where theoretically it is off the equilibrium path.

4 Experimental Findings

Section 4.1 analyzes average behavior. Section 4.2 examines the heterogeneity in individual behavior. Section 4.3 investigates the determinants of individual behavior with regressions.

4.1 Average Behavior

We explore the extent to which average behavior is consistent with the within- and between-treatment comparisons informed by the equilibrium characterizations.

Expert Average Behavior. We use session-level data from the last 20 rounds to capture reasonably converged average behavior. Figure 3 presents the relative frequencies of the five ranking methods. A visual inspection suggests that observed choices vary across ranking value more than across report fee. The frequency distributions are similar in *LL* and *LM*, which differ from those in *HL* and *HH*. There is a more noticeable difference between *HL* and *HH*, suggesting that report fee has a stronger effect when the ranking value is high.

[19] Comparing the average frequencies of methods in the equilibrium and non-equilibrium groups controls for the effect of their uneven sizes given the odd number of ranking methods. Using instead combined frequencies may unfoundedly favor the hypotheses even when choices are haphazard.

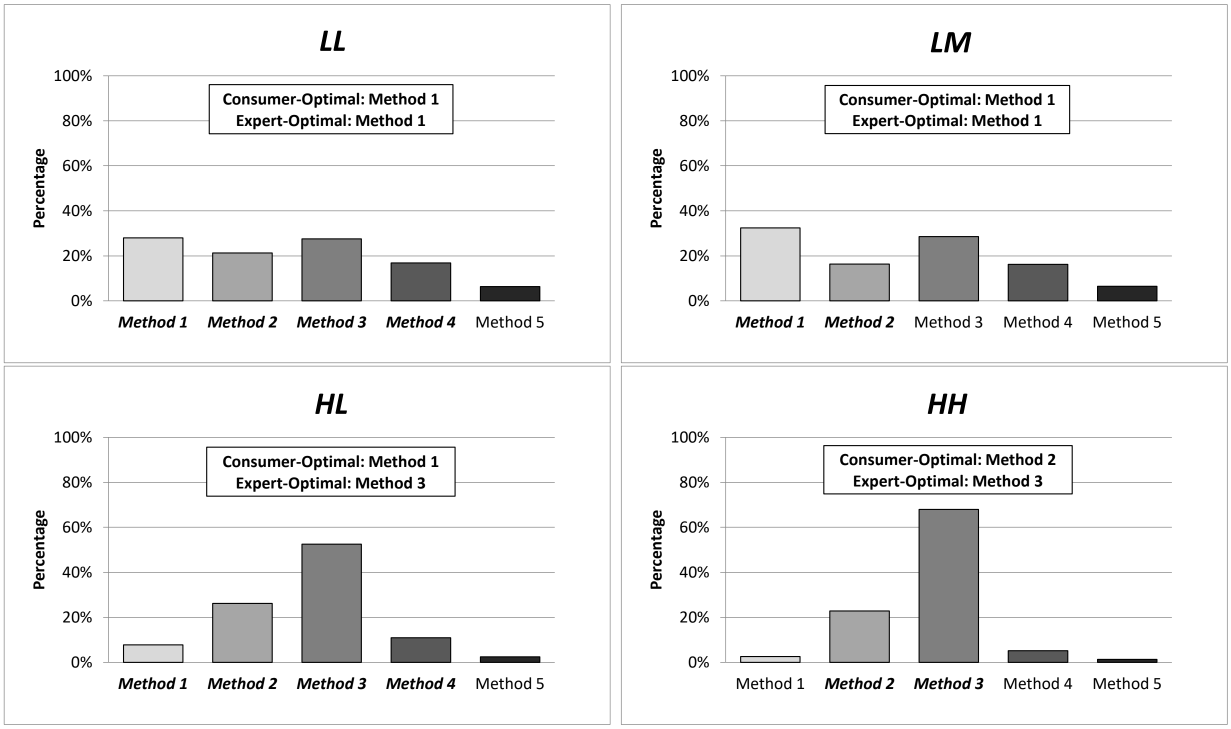

Figure 3: Relative Frequencies of Choices of Ranking Methods

As the first step to evaluate the data in light of theory, we use the perfect Bayesian acquisition equilibria as a gauge to see if average behavior is broadly governed by equilibrium incentives. The acquisition-equilibrium methods are labeled in italicized-bold fonts in Figure 3. Table 5 further consolidates the data into average relative frequencies of equilibrium and non-equilibrium methods. The within-treatment comparisons support the qualitative hypothesis that the equilibrium ranking methods are on average chosen more often than the non-equilibrium methods ($p = 0.0625$ for all treatments, Wilcoxon signed rank tests).[20]

Quantitatively, the finding from *HH* is most remarkable, where the two equilibrium methods together account for 90.8% of the observations. On the other hand, the difference between the equilibrium and non-equilibrium groups, though statistically significant, is smallest in *LM*, where the two equilibrium methods account for 48.8% of the observations.

For the treatment effects of ranking value, acquisition equilibria predict an invariance between *LL* and *HL*: Methods 1–4 are equilibria in both treatments.[21] We evaluate the predicted invariance with a strong and a weak hypotheses. Factors that predict how choices are made amid multiple equilibria are outside the realm of the equilibrium concept. Assuming that these unobserved factors do not systematically change between *LL* and *HL*, we hypothesize that the

[20]Unless otherwise indicated, our non-parametric statistical tests are performed using session-level observations, and reported p-values are from one-sided tests. Note that with four independent observations, $p = 0.0625$ is the *lowest possible* p-value considered as statistically significant for the Wilcoxon signed rank test.

[21]For the treatment effects of ranking value, we do not compare *LM* and *HH* because there is more than one treatment variation, even though theory provides a basis for any pairwise comparison of treatments.

Table 5: Relative Frequencies of Acquisition
Equilibrium and Non-Equilibrium Methods

Treatment	Equilibrium Methods	Non-Equilibrium Methods	Wilcoxon Signed Rank Test
LL	23.4% (4)	6.4% (1)	$p = 0.0625$
LM	24.4% (2)	17.1% (3)	$p = 0.0625$
HL	24.4% (4)	2.5% (1)	$p = 0.0625$
HH	45.4% (2)	3.1% (3)	$p = 0.0625$

Note: A percentage number represents the average relative frequency of the ranking method(s) in the equilibrium or non-equilibrium group. The parenthesis contains the number of methods in the group. The p-values are from one-sided tests ($p = 0.0625$ is the *lowest possible* value with four independent observations).

frequency distributions of the four methods are identical in both treatments. A chi-square test of homogeneity rejects this strong hypothesis ($p < 0.01$).[22]

A weaker hypothesis that uses the dichotomous equilibrium and non-equilibrium groups of methods to gauge the invariance is nevertheless supported: The average relative frequencies of Methods 1–4 and the non-equilibrium Method 5 are 23.4% and 6.4% in *LL*, compared with 24.4% and 2.5% in *HL* (two-sided $p = 0.34$, Mann-Whitney test). While this may merely be a manifestation of the less focal Method 5 being rarely chosen, we contend that the observation is not at odds with the prediction.

Acquisition equilibria predict that the report fee has a narrowing effect: The set of equilibrium ranking methods shrinks from Methods 1–4 in *LL* to Methods 1–2 in *LM* and from Methods 1–4 in *HL* to Methods 2–3 in *HH*. The treatment effect, which we evaluate by looking at methods that are equilibrium in one treatment but not the other, is observed only for high ranking value: The average relative frequency of Methods 1 and 4 is 3.9% in *HH*, significantly lower than the 9.4% in *HL* ($p = 0.03$, Mann-Whitney test). The average relative frequencies of Methods 3 and 4 are, by contrast, virtually the same in *LL* and *LM* (22.4% vs. 22.2%, two-sided $p = 1$, Mann-Whitney test).

We summarize the above findings:

Finding 1. *Evaluating expert average behavior yields the following findings that conform with the predictions of acquisition equilibria:*

(a) The equilibrium ranking methods are chosen more often than the non-equilibrium methods.

(b) A higher ranking value has no effect on the relative frequency distribution over the di-

[22]The chi-square test compares the distributions of the frequency counts of the five methods between the two treatments, rather than the relative frequencies in percentage terms that are reported.

chotomous groups of equilibrium and non-equilibrium methods.

(c) A higher report fee narrows the choices of ranking methods, but it only occurs under high ranking value.

Finding 1 suggests that expert average behavior is overall consistent with the broad predictions of acquisition equilibria. We embark on the more demanding test of the theory, further evaluating the observations in light of the unique predictions of the consumer-optimal and the expert-optimal equilibria. For within-treatment comparisons, we single out the modal choice and juxtapose it with the uniquely predicted method.

The predictions of the consumer-optimal and the expert-optimal equilibria diverge for the treatments with high ranking value. The data favor the latter. The expert-optimal Method 3 is modal and chosen 52.6% and 68% of the time in *HL* and *HH* respectively. In both cases, it is significantly more frequent than the second-place Method 2 ($p = 0.0625$, Wilcoxon signed rank tests). On the other hand, the consumer-optimal Method 1 in *HL* and Method 2 in *HH* are chosen only 7.8% and 22.9% of the time respectively.

The predictions of the two selected equilibria coincide for the treatments with low ranking value. The data shows qualitative support for the common prediction but lacks statistical significance. The consumer/expert-optimal Method 1 is modal, chosen 28.0% of the time in *LL* and 32.4% in *LM*. It is, however, not significantly more frequent than the second-place Method 3 ($p \geq 0.31$, Wilcoxon signed rank tests). Quantitatively, the frequencies are also rather low considering that Method 1 is the common unique prediction. We discuss in Section 4.2 a plausible explanation that combines this departure with observations from consumers.

The expert-optimal equilibria also stand out in between-treatment comparisons. For the treatment effects of ranking value, the consumer-optimal equilibria predict no difference between *LL* and *HL* with Method 1 commonly predicted; the expert-optimal equilibria predict that the higher ranking value shifts the choices from Method 1 in *LL* to Method 3 in *HL*. The data support the shift: The relative frequency of Method 3 is 52.6% in *HL*, significantly higher than the 27.5% in *LL*, and the relative frequency of Method 1 is 7.8% in *HL*, significantly lower than the 28.0% in *LL* ($p = 0.01$ in both cases, Mann-Whitney tests).

For the effects of report fee, the differentiation between the two selected equilibria lies in the treatments with high ranking value (both equilibria predict the same for *LL* and *LM*): The consumer-optimal equilibria predict that the higher report fee shifts choices from Method 1 in *HL* to Method 2 in *HH*, while the expert-optimal equilibria predict invariance with Method 3 being selected in both treatments. The data support the invariance of the expert-optimal equilibria: The relative frequencies of Method 3 are not significantly different between *HL* and *LL* (52.6% vs. 68%, two-sided $p = 0.49$, Mann-Whitney test). For the shift predicted by the consumer-optimal equilibria, one leg is missing: While the relative frequency of Method 1 in *HL* at 7.8% is significantly higher than the 2.6% in *HH* ($p = 0.01$, Mann-Whitney test), those of Method 2 in *HL* and *HH* are not significantly different (26.2% vs. 22.9%, two-sided $p = 0.69$, Mann-Whitney test), overall not supporting the shift.

We summarize the findings comparing the two selected equilibria:

Finding 2. *Average choices are more consistent with the expert-optimal equilibria than the consumer-optimal equilibria. In particular, under high ranking value where their predictions diverge, the expert-optimal methods are most frequently chosen by a considerable margin.*

The consumer-optimal equilibria can be motivated by a small agree of altruism toward the acquiring consumer, while the expert-optimal equilibria reflect a strategic sales motive that is spiteful to the non-acquiring consumer. Finding 2 suggests that the sales motive trumps any altruistic motive. In the treatments with high ranking value, the two motives present a tradeoff: either benefit the acquiring consumers with product guidance but render the report more dispensable, or shuffle to harm the non-acquiring consumers with ranking uncertainty yet make doing without the report less tolerable. The prevalence of the expert-optimal equilibria over the consumer-optimal equilibria suggests that experts are driven more by the sales motive.[23]

Recall that selling the reports earns experts the same rewards in all treatments. The documented responses to the treatment variations are presumably via the endogenous choices of consumers. We turn next to this linkage, examining the average behavior of consumers.

Consumer Average Behavior. Figure 4 shows the relative frequencies of report acquisitions conditional on ranking methods. A pattern comparable to the experts' is observed: The differences across the rows of the figure panels, which capture the effects of ranking value, are overall more pronounced than the column differences, which capture the effects of report fee.

To evaluate the observations in light of theory, note that sequential rationality predicts consumers to always acquire reports under the equilibrium ranking methods (the three equilibrium concepts predict the same). Table 6 consolidates the relative frequencies by whether the methods are equilibrium or not. Consumers on average acquire reports more often when theory predicts that they should than when it predicts that they should not. The within-treatment comparisons are significant in all but one treatment ($p = 0.125$ in *LL* and $p = 0.0625$ in *LM*, *HL*, and *HH*, Wilcoxon signed rank tests).

The unconditional total relative frequencies of report acquisitions provide a handy summary variable to evaluate the treatment effects. The numbers reported in the last column of Table 6

[23] The Bayesian persuasion literature (Kamenica and Gentzkow, 2011) focuses on the sender-optimal equilibria. Our finding about the expert-optimal equilibria in an environment that shares commonalities with Bayesian persuasion hints at an empirical foundation for the sender-optimal equilibria and suggests an avenue for further studies to explore the issue.

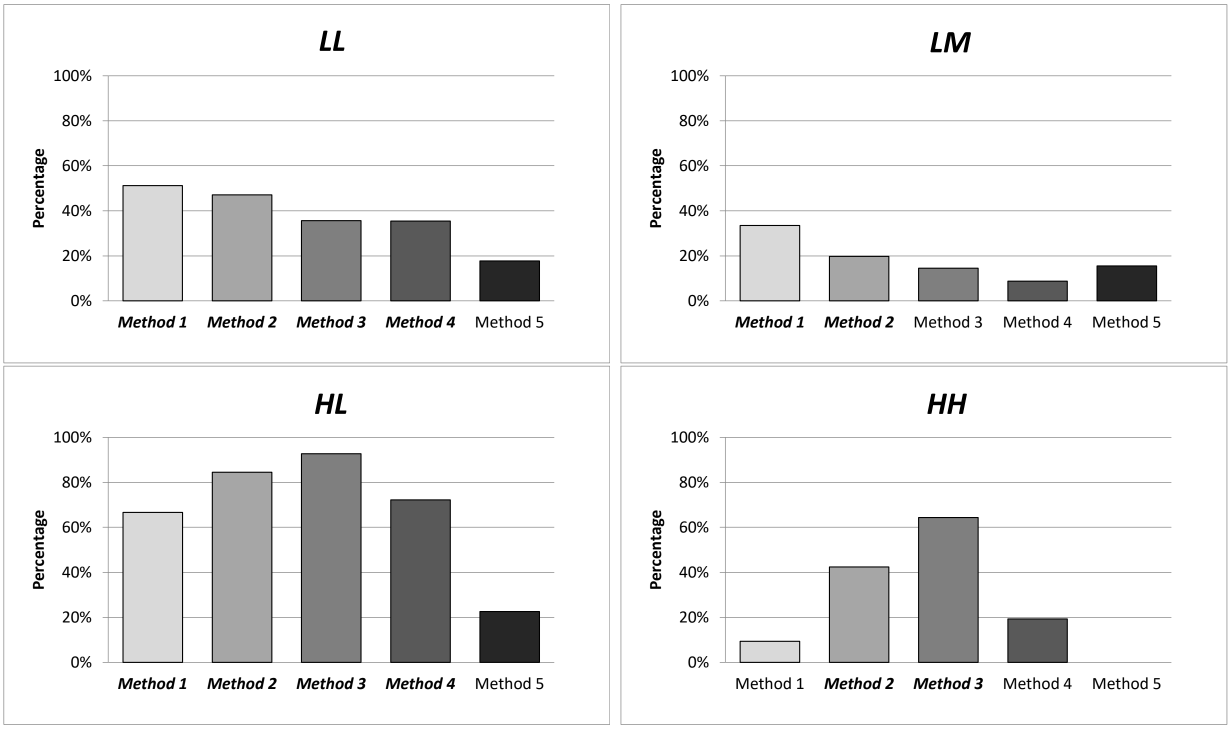

Figure 4: Relative Frequencies of Report Acquisitions

Table 6: Relative Frequencies of Report Acquisitions By Equilibrium and Non-Equilibrium Ranking Methods

Treatment	Equilibrium Methods	Non-Equilibrium Methods	Wilcoxon Signed Rank Test	Total
LL	42.4%	17.7%	$p = 0.125$	42.0%
LM	26.6%	12.9%	$p = 0.0625$	21.7%
HL	79.0%	22.6%	$p = 0.0625$	85.2%
HH	53.4%	10.6%	$p = 0.0625$	55.6%

Note: The p-values are from one-sided tests. With four independent observations, $p = 0.0625$ is the lowest and $p = 0.125$ the second *lowest possible* p-values.

show that acquisition decisions respond to the "benefit-cost ratio" of the reports. Consumers acquire most often when the reports benefit most with high ranking value and cost least with low report fee: In *HL*, reports are acquired 85.2% of the time, which stands in stark contrast to *LM* with the lowest benefit-cost ratio, where the acquisition frequency is only 21.7%.[24] While these are not part of the equilibrium predictions, the treatment effects provide clear evidence that consumers respond to the induced incentives.

We summarize the findings about report acquisitions:

[24]Statistically, the Mann-Whitney tests reveal that the relative frequencies are, compared to *LL* as the baseline, marginally significantly lower in *LM* ($p = 0.06$), significantly higher in *HL* ($p = 0.01$), and not significantly different in *HH* (two-sided $p = 0.34$).

Finding 3. *Ranking reports are acquired more often under equilibrium ranking methods than under non-equilibrium methods. Consumers respond to the treatment incentives, acquiring reports most often when the ranking value and report fee present the highest benefit-cost ratio.*

Consumers' product choices further support the predictions of sequential rationality: Without reports, a product is chosen more often under the ranking methods where it is the optimal default than under the methods where it is not; with reports, the top-ranked products are nearly always chosen. For brevity, we relegate the supporting analysis to Appendix A.3.

Overall, average behavior is qualitatively consistent with the equilibrium predictions. There are, however, sizeable discrepancies between the observations and the *point* predictions. We next move from average to individual behavior, examining individual heterogeneity to shed light on these discrepancies.

4.2 Individual Heterogeneity

Consumer Extreme Behavior. Extreme behavior highlights heterogeneity. Table 7 reports the proportion of consumers in four outermost categories based on their report-acquisition decisions in the 40 rounds. The diverse heterogeneities in different treatments are most sharply exemplified by the contrast between *LM* and *HL*: 47.3% of the consumers in *LM* acquire reports no more than 10% of the time, while about the same proportion in *HL* (42.5%) acquire no less than 90% of the time. Remarkably, 21% of the consumers in *LM* never acquire reports.

Table 7: Proportions of Consumers in Four Non-Exhaustive Categories of Report-Acquisition Frequencies

Treatment	Never	(0 – 10%]	[90 – 100%)	Always
LL	5.3%	7.9%	5.3%	5.3%
LM	21.0%	26.3%	0.0%	0.0%
HL	0.0%	2.5%	30.0%	12.5%
HH	4.8%	2.4%	2.4%	0.0%

Note: A percentage number represents the proportion of consumers whose relative frequencies of acquiring reports in the 40 rounds lie in the category.

The individual findings dissect one of the primary departures from equilibrium predictions observed in average behavior: the infrequent report acquisitions especially under low ranking value. A fraction of consumers, which is particularly sizable in *LM*, apparently find acquiring reports not worthwhile even when equilibrium predicts otherwise. The suboptimal decisions can be rationalized by positing that the expected payoffs from acquiring and not acquiring the report, and therefore the net gain from acquiring $G(\beta_0) - f$, are subject to random perturbations. When $G(\beta_0) - f$ is close to zero so that consumers are making acquisition decisions in vicinity of indifference, the realizations of the perturbations may cause the decisions to regularly land on the suboptimal side.[25]

The above perspective yields a qualitative implication: The instances of suboptimal acquisition decisions are predicted to decrease across treatments as the magnitude of $G(\beta_0) - f$ increases. Table 8 reports the empirical counterparts of $G(\beta_0) - f$ that support this prediction. We compare consumers' equilibrium payoffs with their observed payoffs, averaging their differences across individual consumers and rounds among cases of suboptimal decisions. Matching Table 8 with Table 7 reveals that as this average difference, which is an unrealized net gain, increases from *LM* to *LL* to *HH* to *HL*, the instances of never acquiring reports as a salient indicator of consumers' suboptimal behavior decrease from *HL* to *HH* to *LL* to *LM*.[26]

Table 8: Unrealized Net Gains from
Suboptimal Report-Acquisition Decisions

Treatment	Non-Acquisition	All
LL	25.29	24.67
LM	16.89	18.56
HL	111.47	106.49
HH	32.26	34.34

Note: The "Non-Acquisition" column contains the unrealized net gains (including the report fee) from suboptimal non-acquisitions of reports. The "All" column contains the unrealized net gains from both suboptimal non-acquisitions and acquisitions. The numbers are average across individual consumers and the 40 rounds.

The departure of consumers in turn rationalizes the other departure: the experts' low aggregate frequencies of predicted ranking methods, especially the commonly selected Method 1 for *LL* and *LM*. The equilibrium predictions build on the prospect of the expert selling the report. If some consumers persistently do not acquire reports when theory predicts that they should, then the choices of experts may also be less amenable to the predictions.

Evolution of Expert Behavior. The joint behavior of some non-conforming experts and consumers provides a coherent account for the seemingly haphazard choices of ranking methods in *LL* and *LM*. Facing a fraction of consumers who rarely acquire reports, some experts might be experimenting with different methods over rounds to explore what might sell. The evolution

[25]Random perturbations to payoffs are the foundation of quantal response equilibrium (McKelvey and Palfrey, 1995). While the development of a full model is beyond the scope of this paper, we leverage the idea of payoff perturbations to rationalize the variation in suboptimal acquisition decisions observed across treatments.

[26]The column "Non-Acquisition" in Table 8 reports the unrealized net gains from failing to acquire, while the "All" column covers all suboptimal decisions including failing to pass. Suboptimal non-acquisition is the predominant type of suboptimal decision, accounting for an overwhelming 92% of the cases on average.

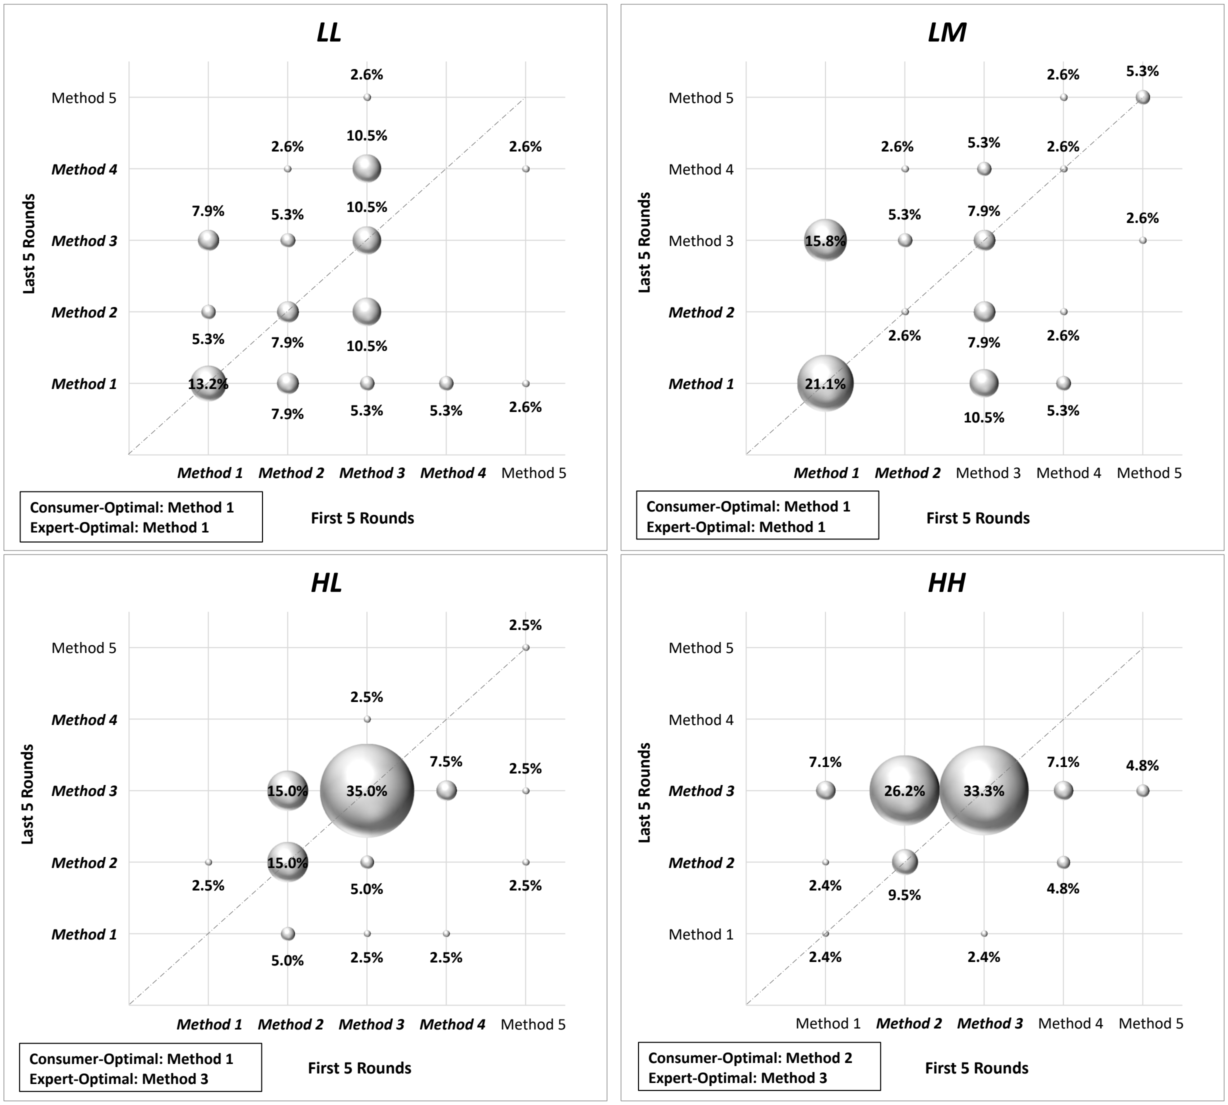

Figure 5: Proportions of Experts in Combinations of Most Frequently Chosen Ranking Methods in First Five and Last Five Rounds

of experts' choices from initial rounds to terminal rounds lends support to this interpretation.

The bubble charts in Figure 5 depict the proportions of experts in different combinations of most frequently chosen methods in the first five and last five rounds.[27] The diagonal in each chart provides an anchor to organize the data in respect of *steady* and *transiting* behavior. Bubbles on the diagonal, which contain experts with coinciding initial and terminal choices, provide a proxy measure of steady behavior. Supporting the thesis of more prevalent experimentations in the low-ranking-value treatments, choices are less steady there than in their high-ranking-value counterparts. The proportions of experts lying on the diagonals are 31.6%

[27] We classify each expert by the ranking methods the expert most frequently chooses in the first and last five rounds (the lowest methods are used in case of ties). A bubble labeled $x\%$ at (Method K, Method K') indicates that $x\%$ of the experts in the treatment have their most frequent choices be Methods K and K' in the initial and terminal rounds respectively. For brevity, we omit "most frequent" in the ensuing discussion.

in *LL* and 39.5% in *LM*, compared to 52.5% in *HL* and 45.2% in *HH*.

The selected equilibrium methods dominate the steady choices. Conditional on those who exhibit steady behavior, 66.7% in *HL* and 73.7% in *HH* choose the expert-optimal Method 3 initially and terminally. Despite the more diffused observations in the low-ranking-value treatments, the consumer/expert-optimal Method 1 accounts for the most common steady behavior there, where the conditional proportions are 41.7% in *LL* and 53.3% in *LM*.

Bubbles off the diagonal outline the patterns of transiting behavior. Experts with non-steady choices in the high-ranking-value treatments commonly gravitate toward the expert-optimal method at the end: Among those with diverging initial and terminal choices, 52.6% in *HL* and 82.6% in *HH* choose Method 3 in the terminal rounds.

For the low-ranking-value treatments, the footprint of the selected equilibria on transiting behavior is less pronounced: 30.8% of the experts in *LL* with non-steady choices choose Method 1 at the end, whereas in *LM* the conditional proportion is 26.1%. The multitude of similarly sized bubbles instead supports the thesis that experts experiment in response to some non-conforming consumers, and they do so rather uniformly across different methods.

Finally, we point out a parallel between the bubble-chart analysis and our main finding on average behavior about the expert-optimal equilibria: Within a treatment, the bubble representing *steady expert-optimal choices* has the greatest sizes among all bubbles; between treatments, the maximal sizes are greater in the high-ranking-value treatments.

We summarize the findings of this subsection:

Finding 4. *Analysis of heterogeneity in individual behavior reveals the following:*

(a) A fraction of consumers more sizable under low ranking value rarely acquire reports.

(b) The proportions of such non-conforming consumers are commensurate with the extents of experimentations by experts over different ranking methods.

(c) Despite the experimentations, steady choices of expert-optimal equilibrium methods in initial and terminal rounds are the modal behavior of individual experts.

4.3 Determinants of Individual Behavior

We further examine the determinants of individual behavior with a regression analysis, taking advantage of the panel data of 158 experts/consumers making decisions in 40 rounds. We estimate binary outcome models using random-effects logit in the following generic form:

$$\Pr(Y_{it} = 1|\mathbf{X}_{it}, \alpha_i) = \Lambda(\mathbf{X}_{it}\boldsymbol{\theta} + \alpha_i), \tag{3}$$

where α_i is the subject-specific effect and $\Lambda(z) = \frac{e^z}{1+e^z}$ is the logistic cumulative distribution.

Determinants of Expert Behavior. For experts, we construct the dependent outcome variable Y_{it} by partitioning the five ranking methods into binary sets. We consider four alternative specifications of Y_{it} that correspond broadly to our four key concepts: product guidance, ranking uncertainty, consumer-optimal equilibria, and expert-optimal equilibria.

For product guidance and ranking uncertainty, it follows from Facts 1 and 2 that Method 1 provides the best guidance and is uncertainty neutral; Method 5 provides the worst guidance and is uncertainty-eliminating; and Method 3 in the middle induces the highest uncertainty.[28] We accordingly specify (a) $Y_{it} = MD_{it}^{\text{wg}}$ and (b) $Y_{it} = MD_{it}^{\text{mu}}$, where MD_{it}^{wg} takes the value of one (zero otherwise) if expert i chooses in round t one of the two methods with *w*orst *g*uidance, Method 4 or 5, and MD_{it}^{mu} takes the value of one (zero otherwise) if expert i chooses in round t one of the two methods that generate the *m*ost *u*ncertain ranking, Method 2 or 3.

For the two selected equilibria, we specify (c) $Y_{it} = MD_{it}^1$ and (d) $Y_{it} = MD_{it}^3$, where MD_{it}^1 takes the value of one (zero otherwise) if expert i chooses in round t Method 1, which is consumer-optimal in three treatments and expert-optimal in two treatments, and MD_{it}^3 is analogously defined for Method 3, which is expert-optimal in two treatments.

For $Y_{it} \in \{MD_{it}^{\text{wg}}, MD_{it}^{\text{mu}}, MD_{it}^1, MD_{it}^3\}$, the specification of $\mathbf{X}_{it}\boldsymbol{\theta}$ on the right-hand side of (3) is given by

$$\mathbf{X}_{it}\boldsymbol{\theta} = \theta_0 + \theta_1 LM_i + \theta_2 HL_i + \theta_3 HH_i + \theta_4 Y_{i,t-1} + \theta_5 SL_{i,t-1} + \theta_6(Y_{i,t-1} \times SL_{i,t-1}),$$

where the independent variables are defined and motivated as follows:

(a) Treatment effects: LM_i, HL_i, and HH_i

- LM_i, HL_i, and HH_i each take the value of one (zero otherwise) if expert i is in the respective treatment.
- θ_1, θ_2, and θ_3 measure how, relative to the baseline LL, the experts in the respective treatments are more or less likely to choose the ranking method(s). For $Y_{it} = MD_{it}^3$, e.g., the expert-optimal equilibria would predict that $\theta_1 = 0$, $\theta_2 > 0$, and $\theta_3 > 0$.

(b) Choice persistence and experience: $Y_{i,t-1}$, $SL_{i,t-1}$, and $Y_{i,t-1} \times SL_{i,t-1}$

- $Y_{i,t-1}$ takes the value of one (zero otherwise) if expert i chooses in round $t-1$ the ranking method(s) in the corresponding case of dependent variable.

[28] In our design, the natural level of ranking uncertainty corresponds to $p = 0.2$, which equals the *ex-ante* probability that Product B is ranked first under Method 1. Product B is always ranked first under Method 5. The *ex-ante* probabilities that the two products are ranked first under Method 3 are the most uniform among the five methods. Note also that Method 4 shares the same ranking uncertainty with Method 1.

- $SL_{i,t-1}$ takes the value of one (zero otherwise) if expert i sells the ranking report in round $t-1$.
- If experts exhibit persistence in their choices irrespective of whether the chosen ranking methods lead to sales, then we will expect that $\theta_4 > 0$.
- If experience matters, then—using $Y_{it} = MD^3_{it}$ as an example—a previous successful experience of selling the report without Method 3 may decrease the odds that the method would be chosen again ($\theta_5 < 0$), and a previous successful selling experience with Method 3 may increase the odds that it would be chosen again ($\theta_6 > 0$).

Columns (1) and (2) of Table 9 report the estimation results for outcomes MD^1_{it} and MD^3_{it} respectively. The coefficients of the treatment dummies indicate that, relative to the baseline *LL*, the experts in *HL* and *HH* are significantly less likely to choose Method 1 and more likely to choose Method 3, while the behavior in *LM* is not that different from the baseline. The significant differences lend support to the expert-optimal equilibria: When a consumer-optimal and expert-optimal method becomes not expert-optimal, it is less likely to be chosen, and when a method is expert-optimal, it is more likely to be chosen even if it is not consumer-optimal.

Table 9: Choices of Ranking Methods: Treatment Effects and Behavioral Determinants

	$Y_{it} = MD^1_{it}$	$Y_{it} = MD^3_{it}$	$Y_{it} = MD^{\text{wg}}_{it}$	$Y_{it} = MD^{\text{mu}}_{it}$
	(1)	(2)	(3)	(4)
LM_i	0.438*	−0.066	−0.122	−0.234
	(0.201)	(0.265)	(0.434)	(0.187)
HL_i	−0.958***	0.692*	−0.329	0.960***
	(0.213)	(0.289)	(0.369)	(0.226)
HH_i	−1.206***	1.243***	−0.628*	1.362***
	(0.371)	(0.254)	(0.303)	(0.151)
$Y_{i,t-1}$	1.243***	0.464**	0.657***	0.693***
	(0.207)	(0.166)	(0.163)	(0.181)
$SL_{i,t-1}$	−0.582**	−0.877***	−0.883***	−1.031***
	(0.189)	(0.161)	(0.164)	(0.218)
$Y_{i,t-1} \times SL_{i,t-1}$	1.934***	1.855***	1.576***	1.932***
	(0.479)	(0.286)	(0.343)	(0.340)
Constant	−2.280***	−1.230***	−1.606***	−0.274
	(0.133)	(0.205)	(0.323)	(0.156)
Observations	6162	6162	6162	6162

Note: Columns (1)–(4) report estimates from four different specifications of dependent variables. The independent variable $Y_{i,t-1}$ is the one-round lagged value of the dependent variable of the column. Standard errors clustered at the session level are in parentheses. *** indicates significance level at 0.1%, ** at 1%, and * at 5%.

The above findings suggest that the experts in the high-ranking-value treatments are more inclined to induce ranking uncertainty with Method 3, which is further substantiated by the result in column (4) where Method 2, which induces the second highest uncertainty, is added to Method 3 for the outcome variable. On product guidance, while the experts in *HL* and *HH* are less inclined to choose Method 1 with the best guidance, the estimates in column (3) indicate that they are not more likely to go all the way to offer the worst guidance with Method 4 or 5.

Note also that for all four outcomes the coefficients of HH_i are greater in magnitudes than those of HL_i. This suggests that the treatment effects are stronger when incentives are more salient under the higher report fee.

Regarding persistence and experience, the significant coefficients with all four outcome variables indicate that experts are prone to repeat their choices, less likely to choose a method when they are able to sell the report in the previous round with another method, and more likely to choose a method that enables them to sell in the previous round. These findings provide further evidence for the experimentations by experts discussed in Section 4.2.

Determinants of Consumer Behavior. For consumers, we use report-acquisition decisions as the outcome variable: $Y_{it} = AQ_{it}$, which takes the value of one (zero otherwise) if consumer i acquires the report in round t. To keep the regression equations as simple as possible, we examine the treatment effects and behavioral determinants separately with two sets of specifications of independent variables. Slightly abusing notation by recycling the use of $\boldsymbol{\theta}$, the first specification to evaluate the treatment effects is

$$\mathbf{X}_{it}\boldsymbol{\theta} = \theta_0 + \theta_1 LM_i + \theta_2 HL_i + \theta_3 HH_i + \theta_4 MD_{it}^1 + \theta_5 MD_{it}^2 + \theta_6 MD_{it}^3 + \theta_7 MD_{it}^4. \tag{4}$$

The regression controls for ranking methods, where MD_{it}^j takes the value of one (zero otherwise) if consumer i encounters Method j in round t. With *LL* being the baseline, the treatment effects previously documented under average behavior would predict that $\theta_1 < 0$, $\theta_2 > 0$, and $\theta_3 = 0$.

Motivated by the prior observation that aggregate report acquisitions reflect the benefit-cost ratio of the report, we also estimate an alternative specification for the treatment effects:

$$\mathbf{X}_{it}\boldsymbol{\theta} = \theta_0 + \theta_1 (R/F)_i + \theta_2 MD_{it}^1 + \theta_3 MD_{it}^2 + \theta_4 MD_{it}^3 + \theta_5 MD_{it}^4, \tag{5}$$

where $(R/F)_i$, which measures the ratio of ranking value to report fee of the treatment in which consumer i makes decisions, supersedes the three treatment dummies in (4).

Column (1) of Table 10 reports the estimates of specification (4). The coefficients of the treatment dummies are in line with the effects extrapolated from the aggregate findings. Column (2) contains the estimates of (5). The coefficient of $(R/F)_i$ corroborates the prior observation that the higher the benefit-cost ratio, the more likely that the report would be acquired.[29]

[29]We reiterate that for the results reported in Table 10 the ranking-method dummies serve as control variables, and we postpone discussing their effects until our analysis of behavioral determinants below.

Table 10: Report Acquisitions: Treatment Effects

	AQ_{it}	
	(1)	(2)
LM_i	−1.775**	–
	(0.591)	–
HL_i	2.403***	–
	(0.542)	–
HH_i	−0.097	–
	(0.376)	–
$(R/F)_i$	–	0.068***
	–	(0.012)
MD^1_{it}	2.865***	2.862***
	(0.492)	(0.494)
MD^2_{it}	2.668***	2.682***
	(0.604)	(0.606)
MD^3_{it}	3.145***	3.166***
	(0.679)	(0.682)
MD^4_{it}	1.778***	1.785***
	(0.532)	(0.534)
Constant	−2.923***	−3.881***
	(0.544)	(0.610)
Observations	6320	6320

Note: Standard errors clustered at the session level are in parentheses. *** indicates significance level at 0.1%, ** at 1%, and * at 5%.

The effective use of $(R/F)_i$ as a proxy for the treatment dummies suggests that we may *pierce the veil of the treatment labels*—viewing consumers as simply responding to the induced incentives of ranking values and report fees. This perspective serves as our starting point to further investigate their behavioral determinants.

Thereby bypassing the treatment dummies, our next regression extends on (5) by adding three independent variables that capture persistence and experience: $AQ_{i,t-1}$, $TP_{i,t-1}$, and their interaction, where $TP_{i,t-1}$ not hitherto defined takes the value of one (zero otherwise) if consumer i chooses the top-ranked product in round $t-1$.[30] To better understand how the incentives of ranking values and report fees may interact with the encountered ranking methods, we also include interaction terms between $(R/F)_i$ and MD^j_{it}.

Persistence would predict the coefficient of $AQ_{i,t-1}$ to be positive, and previous experience of choosing the top-ranked product, if matters, would predict the coefficient of $TP_{i,t-1}$ to be

[30]The variable $AQ_{i,t-1}$ captures the same behavior as $SL_{i,t-1}$ used for experts. They, however, differ by the meaning of the index i, where $AQ_{i,t-1}$ captures the previous-round report acquisition of consumer i and $SL_{i,t-1}$ captures the report acquisition of the consumer matched with expert i in the previous round.

Table 11: Report Acquisitions: Behavioral Determinants

	AQ_{it}			
	(1)	(2)	(3)	(4)
$(R/F)_i$	0.009	–	–	–
	(0.018)	–	–	–
$G_{it} - F_i$	–	0.043***	–	0.047***
	–	(0.005)	–	(0.007)
$AQ_{i,t-1}$	−0.059	−0.043	−0.112	−0.025
	(0.260)	(0.279)	(0.267)	(0.263)
$TP_{i,t-1}$	−0.093	−0.075	−0.097	−0.077
	(0.113)	(0.124)	(0.114)	(0.122)
$AQ_{i,t-1} \times TP_{i,t-1}$	0.628*	0.557	0.668*	0.554
	(0.274)	(0.297)	(0.284)	(0.284)
MD^1_{it}	2.165***	–	2.945***	−0.045
	(0.461)	–	(0.533)	(0.388)
MD^2_{it}	1.467**	–	2.729***	−0.071
	(0.489)	–	(0.649)	(0.316)
MD^3_{it}	1.958**	–	3.180***	−0.239
	(0.705)	–	(0.733)	(0.282)
MD^4_{it}	0.519	–	1.847***	−0.209
	(0.362)	–	(0.580)	(0.215)
$(R/F)_i \times MD^1_{it}$	0.023	–	–	–
	(0.015)	–	–	–
$(R/F)_i \times MD^2_{it}$	0.066*	–	–	–
	(0.027)	–	–	–
$(R/F)_i \times MD^3_{it}$	0.065*	–	–	–
	(0.026)	–	–	–
$(R/F)_i \times MD^4_{it}$	0.068***	–	–	–
	(0.014)	–	–	–
LM_i	–	–	−1.674**	−0.575
	–	–	(0.554)	(0.578)
HL_i	–	–	2.212***	−1.174
	–	–	(0.507)	(0.734)
HH_i	–	–	−0.105	0.706
	–	–	(0.335)	(0.386)
Constant	−2.921***	−1.615***	−3.126***	−1.451***
	(0.430)	(0.225)	(0.549)	(0.370)
Observations	6162	6162	6162	6162

Note: Standard errors clustered at the session level are in parentheses. *** indicates significance level at 0.1%, ** at 1%, and * at 5%.

negative and that of $AQ_{i,t-1} \times TP_{i,t-1}$ to be positive. Column (1) of Table 11 reports the estimation result. Unlike experts, consumers are not persistent in their report acquisitions. Neither does stumbling on the top-ranked products without reports have any significant impact. The decisions to acquire are nevertheless moderately reinforced by previous experience of securing the top-ranked products after viewing reports.

The stand-alone effect of $(R/F)_i$ becomes insignificant in the richer specification, but its interactions with the ranking-method dummies points to the sequential rationality exhibited by consumers. Relative to the baseline of Method 5, consumers are significantly more likely to acquire reports under Methods 1, 2, and 3, but not under Method 4; yet the interaction between $(R/F)_i$ and MD^4_{it} is positive and highly significant. This aligns with the fact that acquiring the report under Method 4 is sequentially rational only in *HL* and *LL*, and, among the four treatments, *HL* has the highest $(R/F)_i$ followed by *LL*. On the other hand, the interaction is not significant for Method 1, which is consistent with the fact that it is sequentially rational to acquire the report under Method 1 even in *LM* with the lowest $(R/F)_i$.

The above finding motivates us to go one step further to subsume also the ranking methods into incentive values in our last regression. We construct a summary incentive variable measuring the net expected gain from acquiring the report given the ranking method consumer i encounters in round t: $G_{it} - F_i$, which corresponds to $G(\beta_0) - f$ in Figure 1. The variable supersedes $(R/F)_i$ and MD^j_{it}. The persistence and experience variables are preserved.

Column (2) of Table 11 reports the estimation result. Ranking reports are significantly more likely to be acquired when the net gains are higher. To underscore the significance of this finding, we estimate two additional regressions, in one specification replacing $G_{it} - F_i$ back with the treatment and ranking-method dummies and in the other adding these dummies while keeping $G_{it} - F_i$. Juxtaposing columns (3) and (4) reveals that the significant effects of treatments and ranking methods seen in column (3) vanish once we control for $G_{it} - F_i$. These findings provide the most direct evidence that consumers respond to the induced incentives. We also note that there are no drastic differences in the effects of persistence and experience across the different specifications.

We summarize the key findings of this subsection, which concludes our data analysis:

Finding 5. *Analyses of subject-level panel data corroborate the dominance of expert-optimal equilibria, reveal that consumers respond to the bare incentives behind ranking methods, and show that experts are choice-persistent and influenced by experience, more so than consumers.*

5 Concluding Remarks

Motivated by the lack of structured evidence on the sentiment among industry observers that ranking publishers excessively alter their product rankings for marketing purposes, this study resorts to experimental evidence. Guided by the formal analysis of a ranking-report game that helps make precise the layman view, we use monetary payments to induce in the laboratory plausible incentives faced by ranking publishers.

In our game, "altering the product rankings" manifests as strategic shuffling, in which the expert sometimes ranks the less intrinsically valuable product at the top, even when doing so means not offering product guidance to the consumer. This strategic move engenders a willingness to pay for the expert's ranking report to resolve the uncertainty over which product carries the top-ranked prestige valued by the consumer. Our equilibrium analysis provides a sense that this shuffling, while benefiting the expert, may be done excessively from the vantage point of consumer welfare: When the prestige value of ranking is relatively high, the expert-optimal equilibrium diverges from the consumer-optimal equilibrium.

Our experimental findings show that this excessive shuffling is not only an equilibrium phenomenon but also a laboratory one. Equilibrium ranking methods that are expert-optimal and consumer-optimal, when the two coincide, are chosen more often than other equilibrium methods; more importantly, when they diverge, the expert-optimal method, also the shuffling method inducing ranking uncertainty, is most frequently chosen by a considerable margin. Our experiment provides evidence supporting the view that a profit-driven ranking publisher may adopt ranking methodologies for marketing purposes at the expense of consumers.

We discuss two directions for future research. In our game, the sellers of the products being ranked are not part of the environment. In practice, sellers subscribe to consulting services offered by ranking publishers to stay up-to-date about their ranking criteria, and this may create an additional sales motive for publishers to alter their criteria. Sellers may also play a strategic role in the designs, impacts, and reputations of product rankings. Luca and Smith (2015) provide empirical evidence that business schools selectively promote the publications that favorably rank their programs; in arguably the biggest challenge to the industry, prominent law schools withdrew from the *U.S. News & World Report* law school rankings over objections to the publication's ranking methodology (Gerken, 2022; Manning, 2022), with *U.S. News* changing its methodology to reflect these concerns (Morse and Salmon, 2023). How in a two-sided market sellers and consumers interact with publishers to shape the ranking outcomes represents an important question to be explored.

For a simple experimental environment, we have considered a game without competition, whereas the product-ranking industries are typically made up of multiple publishers ranking the same class of products (other than *Kelley Blue Book*, *Car and Driver* also offers their editor's choices of cars). Competition may drive methodology specialization. In the rankings of undergraduate programs, e.g., one publisher may emphasize student qualities, while another may focus on the value-added of the educational experiences. Moreover, the options to access multiple rankings might dilute the prestige effect from any given ranking. If competition lowers the importance of ranking prestige relative to product guidance, then it may attenuate the consumer-unfriendly incentives to shuffle, echoing the familiar theme that competition benefits consumers. We leave the evaluations of these conjectures to future research.

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, OSub Kwon, and King King Li (2023), "A simple experiment on simple Bayesian persuasion." Working Paper.

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 (2023), "Mediated talk: An experiment." *Journal of Economic Theory*, 208, 105593.

Camerer, Colin (1995), "Individual decision making." In *Handbook of Experimental Economics* (John H. Kagel and Alvin E. Roth, eds.), 587–683, Princeton University Press.

Chakraborty, Archishman and Rick Harbaugh (2007), "Comparative cheap talk." *Journal of Economic Theory*, 132, 70–94.

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.

Chen, Roy and Yan Chen (2011), "The potential of social identity for equilibrium selection." *American Economic Review*, 101, 2562–2589.

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*, 50, 1431–1451.

de Groot Ruiz, Adrian, Theo Offerman, and Sander Onderstal (2015), "Equilibrium selection in experimental cheap talk games." *Games and Economic Behavior*, 91, 14–25.

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.

Fréchette, Guillaume, Alessandro Lizzeri, and Jacopo Perego (2022), "Rules and commitment in communication: An experimental analysis." *Econometrica*, 90, 2283–2318.

Fudenberg, Drew and Emanuel Vespa (2019), "Learning theory and heterogeneous play in a signaling-game experiment." *American Economic Journal: Microeconomics*, 11, 186–215.

Gerken, Heather (2022), "Why Yale law school is leaving the U.S. News & World Report rankings." *Yale Law School*. https://law.yale.edu/yls-today/news/dean-gerken-why-yale-law-school-leaving-us-news-world-report-rankings.

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.

Hagenbach, Jeanne and Eduardo Perez-Richet (2018), "Communication with evidence in the lab." *Games and Economic Behavior*, 112, 139–165.

Jin, Ginger Zhe, Michael Luca, and Danel Martin (2021), "Is no news (perceived as) bad news? An experimental investigation of information disclosure." *American Economic Journal: Microeconomics*, 13, 141–73.

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

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.

Lai, Ernest K. and Wooyoung Lim (2018), "Meaning and credibility in experimental cheap-talk games." *Quantitative Economics*, 9, 1453–1487.

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.

Manning, John (2022), "Decision to withdraw from the U.S. News & World Report process." *Harvard Law Today*. https://hls.harvard.edu/today/decision-to-withdraw-from-the-u-s-news-world-report-process/.

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

McKelvey, Richard D. and Thomas R. Palfrey (1995), "Quantal response equilibria for normal form games." *Games and Economic Behavior*, 10, 6–38.

Milgrom, Paul and Joshua Mollner (2021), "Extended proper equilibrium." *Journal of Economic Theory*, 194, 105258.

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

Morse, Robert and Stephanie Salmon (2023), "Plans for publication of the 2023–2024 best law schools." *U.S. News*. https://www.usnews.com/education/blogs/college-rankings-blog/articles/2023-01-02/plans-for-publication-of-the-2023-2024-best-law-schools.

Myerson, Roger (1978), "Refinements of the nash equilibrium concept." *International Journal of Game Theory*, 7, 73–80.

Ottaviani, Marco and Peter Norman Sørensen (2006), "Reputational cheap talk." *Rand Journal of Economics*, 37, 155–175.

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.

Schmidt, William and Ryan W. Buell (2017), "Experimental evidence of pooling outcomes under information asymmetry." *Management Science*, 63, 1271–1656.

Shannon, Claude (1948), "A mathematical theory of communication." *Bell System Technical Journal*, 27, 379–423, 623–656.

Simon, Leo K. and Maxwell B. Stinchcombe (1995), "Equilibrium refinement for infinite normal-form games." *Econometrica*, 63, 1421–1443.

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 U.S. News & World Report college rankings." *The Atlantic*. https://www.theatlantic.com/education/archive/2013/09/your-annual-reminder-to-ignore-the-em-us-news-world-report-em-college-rankings/279103/.

Wu, Wenhao and Bohan Ye (2023), "Competition in persuasion: An experiment." *Games and Economic Behavior*, 138, 72–89.

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 and Additional Analysis

A.1 Proofs and Verifications

Proof of Lemma 1. Given β_0 and Assumption 2, Product A is the optimal default product if and only if $V_A^{\varnothing}(\beta_0) = \bar{v}_A + (1-p)(1-\beta_0)r$ is greater than or equal to $V_B^{\varnothing}(\beta_0) = p\bar{v}_B + [p + (1-p)\beta_0]r$. Since $\frac{\partial[V_A^{\varnothing}(\beta_0) - V_B^{\varnothing}(\beta_0)]}{\partial \beta_0} = -2(1-p)r < 0$, $V_A^{\varnothing}(\beta_0) - V_B^{\varnothing}(\beta_0)$ is strictly decreasing in β_0. It follows that Product A is the optimal default if and only if $\beta_0 \le \beta_{AB} = \frac{\bar{v}_A - p\bar{v}_B + (1-2p)r}{2(1-p)r}$.

□

Proof of Proposition 1. Given β_0 and Assumption 2, the consumer acquires the ranking report if and only if $G(\beta_0) \ge f$. The equilibrium choices of β_0 then follow from the fact that the expert strictly prefers to sell the report, is indifferent between any β_0 under which the report will be sold, and is indifferent between any β_0 under which the report will not be sold.

□

Verification of Cases in Table 1. We verify the cases by considering the different scenarios of $G(\beta_0) \ge f$ and characterizing the consumer's equilibrium (sequentially rational) strategies. The consumer prefers acquiring the report over not acquiring with Product A as the default if and only if $V(\beta_0) - V_A^{\varnothing}(\beta_0) - f \ge 0$, where $V(\beta_0) = p\bar{v}_B + (1-p)(1-\beta_0)\bar{v}_A + r$ and $V_A^{\varnothing}(\beta_0) = \bar{v}_A + (1-p)(1-\beta_0)r$. This is equivalent to

$$\beta_0 \le \beta_A = \frac{f - f_1}{(1-p)(r - \bar{v}_A)} \text{ and } r < \bar{v}_A, \text{ or} \tag{6}$$

$$\beta_0 \ge \beta_A = \frac{f - f_1}{(1-p)(r - \bar{v}_A)} \text{ and } r > \bar{v}_A, \tag{7}$$

where $f_1 = p(\bar{v}_B + r - \bar{v}_A)$. The consumer prefers acquiring the report over not acquiring with Product B as the default if and only if $V(\beta_0) - V_B^{\varnothing}(\beta_0) - f \ge 0$, where $V_B^{\varnothing}(\beta_0) = p\bar{v}_B + [p + (1-p)\beta_0]r$. This is equivalent to

$$\beta_0 \le \beta^B = 1 - \frac{f}{(1-p)(\bar{v}_A + r)}. \tag{8}$$

The preference conditions (6), (7), and (8) are established by fixing a default product as alternative to acquiring the report. Sequential rationality requires the default alternative to be optimal, and the condition from Lemma 1 for Product A to be the optimal default is

$$\beta_0 \le \beta_{AB} = \frac{\bar{v}_A - p\bar{v}_B + (1-2p)r}{2(1-p)r}, \tag{9}$$

where the four cases are based on parameters that satisfy $0 \le \beta_{AB} < 1$. We let $f_2 = \frac{(p\bar{v}_B - \bar{v}_A + r)(\bar{v}_A + r)}{2r}$ and use conditions (6)–(9) and Assumption 2 to complete the verification.

We use $a_{K'}^{K}$ to denote the product choice rule with the property that Product $K \in \{A, B, TP\}$ is chosen upon viewing the ranking report, where TP denotes the top-ranked product, and Product $K' \in \{A, B\}$ is the default product.

For Cases 1 and 2 where $r < \bar{v}_A$ so that $f_2 \le f_1$, if $f \in (0, f_2]$, then the three thresholds in (6), (8), and (9) satisfy $0 \le \beta_{AB} \le \beta_B \le \beta_A$, and if $f \in (f_2, f_1]$, then the thresholds satisfy $0 \le \beta_A < \beta_B < \beta_{AB} < 1$. The following thus constitute the sequentially rational strategies of the consumer: For Case 1, $s(\beta_0) = 1$ with a_A^{TP} for $\beta_0 \in [0, \beta_{AB}]$, $s(\beta_0) = 1$ with a_B^{TP} for $\beta_0 \in (\beta_{AB}, \beta_B]$, and $s(\beta_0) = 0$ with a_B^{TP} for $\beta_0 \in (\beta_B, 1]$. For Case 2, $s(\beta_0) = 1$ with a_A^{TP} for $\beta_0 \in [0, \beta_A]$, $s(\beta_0) = 0$ with a_A^{TP} for $\beta_0 \in (\beta_A, \beta_{AB}]$, and $s(\beta_0) = 0$ with a_B^{TP} for $\beta_0 \in (\beta_{AB}, 1]$.

For Cases 3 and 4 where $r > \bar{v}_A$ so that $f_1 \le f_2$, if $f \in (0, f_1]$, then the three thresholds in (7), (8), and (9) satisfy $\beta_A \le 0 \le \beta_{AB} \le \beta_B < 1$, and if $f \in (f_1, f_2]$, then the thresholds satisfy $0 < \beta_A \le \beta_{AB} \le \beta_B < 1$. The following thus constitute the sequentially rational strategies of the consumer: For Case 3, $s(\beta_0) = 1$ with a_A^{TP} for $\beta_0 \in [0, \beta_{AB}]$, $s(\beta_0) = 1$ with a_B^{TP} for $\beta_0 \in (\beta_{AB}, \beta_B]$, and $s(\beta_0) = 0$ with a_B^{TP} for $\beta_0 \in (\beta_B, 1]$. For Case 4, $s(\beta_0) = 0$ with a_A^{TP} for $\beta_0 \in (0, \beta_A)$, $s(\beta_0) = 1$ with a_A^{TP} for $\beta_0 \in [\beta_A, \beta_{AB}]$, $s(\beta_0) = 1$ with a_B^{TP} for $\beta_0 \in (\beta_{AB}, \beta_B]$, and $s(\beta_0) = 0$ with a_B^{TP} for $\beta_0 \in (\beta_B, 1]$.

□

Verification of Cases in Table 2. We verify the cases by solving the corresponding maximization problems. For the consumer-optimal equilibrium, since $\frac{\partial V(\beta_0)}{\partial \beta_0} = -(1-p)\bar{v}_A < 0$, the unique solution to the maximization problem is the lower bound of the equilibrium range in each of the four cases.

For the expert-optimal equilibrium, there are two instances for the derivative of $G(\beta_0)$: (i) $\frac{\partial G(\beta_0)}{\partial \beta_0} = (1-p)(r-\bar{v}_A)$ for the instance where $V_A^{\varnothing}(\beta_0) \ge V_B^{\varnothing}(\beta_0)$, and (ii) $\frac{\partial G(\beta_0)}{\partial \beta_0} = -(1-p)(\bar{v}_A + r)$ for the instance where $V_A^{\varnothing}(\beta_0) < V_B^{\varnothing}(\beta_0)$. For Cases 1 and 2 where $r < \bar{v}_A$, $\frac{\partial G(\beta_0)}{\partial \beta_0} < 0$ in both instances (i) and (ii), and thus the unique solution to the maximization problem is $\beta_0 = 0$ in each of Cases 1 and 2. For Cases 3 and 4 where $r > \bar{v}_A$, $\frac{\partial G(\beta_0)}{\partial \beta_0} > 0$ in instance (i) and $\frac{\partial G(\beta_0)}{\partial \beta_0} < 0$ in instance (ii). From the equilibrium strategies of the consumer in Cases 3 and 4 spelled out in the verification for Table 1, the unique relevant solution to the maximization problem is $\beta_0 = \beta_{AB}$ obtained from instance (i).

□

A.2 Selecting Expert-Optimal Equilibrium

We develop a perturbation-based refinement to select the expert-optimal equilibrium. Though *ad hoc*, devised specifically for our game, the refinement that we term *robust acquisition equilibrium* shares the spirit of Myerson's (1978) proper equilibrium: Trembles are introduced to the consumer's report acquisition in such a way that more costly mistakes are less likely.

We denote by $s(\beta_0)$ the consumer's optimal report-acquisition decision when the expert chooses β_0 and by $a(\beta_0, s)$ her optimal product choice when the expert chooses β_0 and her report-acquisition decision is s. For $s \neq s(\beta_0)$, we define $G(\beta_0, s) = V(\beta_0, s(\beta_0), a(\beta_0, s(\beta_0))) - V(\beta_0, s, a(\beta_0, s))$, where V equals the expected utility in either (1) or (2). Then, $\mathcal{G}(\beta_0, s) = G(\beta_0, s) - (1 - 2s)f$ is the consumer's expected payoff gain from choosing the optimal $s(\beta_0)$ instead of the non-optimal $s \neq s(\beta_0)$ taking into account the effect of the report fee. The consumer's trembles are captured by a mixed report-acquisition rule $\sigma : [0,1] \times \{0,1\} \to [0,1]$, where $\sigma(\beta_0, s)$ is the probability that she chooses $s \in \{0,1\}$ given β_0.

For $\epsilon > 0$, we define a "mistake function" $e_\epsilon : [0,1] \times \{0,1\} \to (0, \epsilon)$, where $e_\epsilon(\beta_0, s)$ is the minimum weight the mixed report-acquisition rule puts on s in the subgame set off by β_0. A mistake function satisfies *strict loss monotonicity* if for all $\tilde{\beta}_0, \hat{\beta}_0 \in [0,1]$ and all $s', s'' \in \{0,1\}$, $\mathcal{G}(\tilde{\beta}_0, s') > \mathcal{G}(\hat{\beta}_0, s'')$ implies that $e_\epsilon(\tilde{\beta}_0, s') < e_\epsilon(\hat{\beta}_0, s'')$. The definition of robust acquisition equilibrium leverages the restriction of this monotonicity:[31]

Definition 1 (Robust Acquisition Equilibrium)**.** *For $\epsilon > 0$, a strategy profile $(\beta_0, (\sigma_\epsilon, a))$ with totally mixed σ_ϵ is an ϵ-constrained acquisition equilibrium if*

(a) β_0 is the expert's optimal choice of ranking method given (σ_ϵ, a),

(b) σ_ϵ is the consumer's constrained optimal report-acquisition rule subject to $\sigma_\epsilon(\beta_0, s) \geq e_\epsilon(\beta_0, s)$ for all $\beta_0 \in [0,1]$, all $s \in \{0,1\}$, and any $e_\epsilon : [0,1] \times \{0,1\} \to (0, \epsilon)$ that satisfies strict loss monotonicity, and

(c) $a = a(\beta_0, s)$ is the consumer's optimal product choice rule.

A robust acquisition equilibrium is any limit of ϵ-constrained acquisition equilibria as $\epsilon \to 0$.

Applying Definition 1, we obtain the following characterization:

[31] Myerson's (1978) proper equilibrium is for finite games and cannot be directly extended to infinite games as there may be uncountably many successively costlier mistakes creating cardinality issues. Simon and Stinchcombe (1995) introduce various adaptations. While we can follow their approaches, we introduce instead the mistake function and strict loss monotonicity for an intuitive construction. Note also that our tremble restrictions can be viewed as being imposed across different agents of the consumer each playing a subgame. See Milgrom and Mollner (2021) for a refinement of proper equilibrium by adding across-player tremble restrictions.

Proposition 2. *If $\hat{\beta}_0$ constitutes a robust acquisition equilibrium, then $\hat{\beta}_0$ maximizes $\mathcal{G}(\beta_0, 0)$ over the set of all acquisition equilibrium β_0.*

Proof of Proposition 2. In any ϵ-constrained acquisition equilibrium, the consumer's totally mixed σ_ϵ satisfies: For any $\beta_0 \in [0,1]$, if $V(\beta_0, 1, a(\beta_0, 1)) - f \geq V(\beta_0, 0, a(\beta_0, 0))$, then $\sigma_\epsilon(\beta_0, 0) = e_\epsilon(\beta_0, 0)$, and if $V(\beta_0, 1, a(\beta_0, 1)) - f < V(\beta_0, 0, a(\beta_0, 0))$, then $\sigma_\epsilon(\beta_0, 1) = e_\epsilon(\beta_0, 1)$, where $e_\epsilon \in (0, \epsilon)$. It follows from the definition of $\mathcal{G}$ that $\sigma_\epsilon(\beta_0, 0) = e_\epsilon(\beta_0, 0)$ for $\beta_0 \in [0,1]$ such that $\mathcal{G}(\beta_0, 0) \geq 0$ and $\sigma_\epsilon(\beta_0, 1) = e_\epsilon(\beta_0, 1)$ for $\beta_0 \in [0,1]$ such that $\mathcal{G}(\beta_0, 1) > 0$. This totally mixed report-acquisition rule induces the following expected payoff for the expert from choosing $\beta_0 \in [0,1]$:

$$\pi\{[1 - e_\epsilon(\beta_0, 0)]\mathbb{I}_{\mathcal{G}} + e_\epsilon(\beta_0, 1)(1 - \mathbb{I}_{\mathcal{G}})\}, \tag{10}$$

where $\mathbb{I}_{\mathcal{G}} \in \{0,1\}$ takes the value of one if $\mathcal{G}(\beta_0, 0) \geq 0$ and zero if $\mathcal{G}(\beta_0, 1) > 0$. Since $1 - e_\epsilon(\beta_0, 0) > e_\epsilon(\beta_0, 1)$ for $e_\epsilon(\beta_0, 0) < \frac{1}{2}$ and $e_\epsilon(\beta_0, 1) < \frac{1}{2}$, in any ϵ-constrained acquisition equilibrium with $\epsilon < \frac{1}{2}$, the expert must choose a β_0 for which $\mathcal{G}(\beta_0, 0) \geq 0$, and the expression of the expert's expected payoff in (10) reduces to $\pi[1 - e_\epsilon(\beta_0, 0)]$.

The strict loss monotonicity of e_ϵ implies that $e_\epsilon(\beta_0, 0)$ is strictly decreasing in $\mathcal{G}(\beta_0, 0)$. Therefore, in any ϵ-constrained acquisition equilibrium with $\epsilon < \frac{1}{2}$, by best responding to the consumer's constrained optimal σ_ϵ, choosing a β_0 that maximizes $\pi[1 - \sigma_\epsilon(\beta_0, 0)] = \pi[1 - e_\epsilon(\beta_0, 0)]$, the expert also chooses a β_0 that maximizes $\mathcal{G}(\beta_0, 0)$. The claim follows by noting that a robust acquisition equilibrium is any limit of ϵ-constrained acquisition equilibria as $\epsilon \to 0$, and thus $\epsilon < \frac{1}{2}$ must hold approaching the limit.

□

By definition the expert-optimal equilibrium β_0 maximizes $\mathcal{G}(\beta_0, 0)$. Furthermore, while Proposition 2 establishes a necessary condition, the unique solution to the maximization implies that a robust acquisition equilibrium exists and is unique in each case. It follows that the robust acquisition equilibrium method coincides with the expert-optimal equilibrium method.

A.3 Data Analysis on Product Choices

Table A.1 presents the relative frequencies of Product B being chosen as the default under two groups of ranking methods. The within-treatment comparisons are qualitatively consistent with the predictions of sequential rationality where Product B is never the optimal default under Methods 1–3.[32] The relative frequency of default Product B is significantly lower under Methods 1 – 3 than under Methods 4 – 5 in all four treatments ($p = 0.0625$, Wilcoxon signed

[32] For *HL* and *HH*, the consumer is indifferent between the two products under Method 3, which is resolved in favor of Product A by Assumption 2.

rank tests). The magnitudes of the differences are greater under high ranking value, and the largest difference is recorded in *HL* with 11.8% vs. 96.9%.

Table A.1: Relative Frequencies of Product Choices

	Without Acquiring Report: Default Product B			Acquiring Report: Top-Ranked Product
Treatment	Methods 1 – 3	Methods 4 – 5	Wilcoxon Signed Rank Test	All Methods
LL	6.5%	40.0%	$p = 0.0625$	91.9%
LM	8.7%	44.0%	$p = 0.0625$	95.2%
HL	11.8%	96.9%	$p = 0.0625$	98.1%
HH	20.0%	100.0%	$p = 0.0625$	99.3%

Note: The p-values are from one-sided tests. With four independent observations, $p = 0.0625$ is the *lowest possible* p-values for the Wilcoxon signed rank test.

Table A.1 also provides the relative frequencies of the top-ranked products conditional on the acquisitions of reports. The top-ranked products are chosen more than 90% of the time in all treatments. Predictably, the ranking reports influence consumers' product choices.