PRODUCTION AND OPERATIONS MANAGEMENT

PRODUCTION AND OPERATIONS MANAGEMENT

Vol. 27, No. 11, November 2018, pp. 1960–1977
ISSN 1059-1478 | EISSN 1937-5956 | 18 | 2711 | 1960
DOI 10.1111/poms.12592

Online Reviews and Collaborative Service Provision: A Signal-Jamming Model

Haoying Sun
Department of Marketing & Supply Chain, Gatton College of Business and Economics, University of Kentucky, 550 South Limestone, Lexington, Kentucky 40506, USA, haoying.sun@uky.edu

Lizhen Xu
Scheller College of Business, Georgia Institute of Technology, Atlanta, Georgia 30308, USA, lizhen.xu@scheller.gatech.edu

We study the provision of collaborative services under online reviews, where the service outcome depends on the effort of both the service provider and the client. The provider decides not only her own effort but also the client's, at least to some extent. The client gives the review based on his net utility upon service completion. We develop a signal-jamming model in which the provider's inherent capability or type is unobservable, and the market infers the provider type through observable signals such as the service outcome, the client review, or both. We show that compared to the benchmark case when the service outcome is observed as a signal, the client review generally leads to less effort of both the provider and the client. The review hence tends to sacrifice the service effectiveness in favor of the efficiency of the client's effort input. Nevertheless, when clients incorporate private information about the provider type into their reviews, service providers are better motivated to devote effort. Interestingly, we find the provider's effort choices may be either strategic complements or substitutes. With a reasonable level of informativeness, online reviews could lead to favorable performance in service effectiveness, client effort efficiency, and provider type distinguishability. Surprisingly, we demonstrate that when both the review and the outcome are available, the provider may lack sufficient incentive to devote effort, resulting in inferior distinguishability of provider types. It thus illustrates that richer information may not necessarily generate favorable strategic outcomes.

Key words: online review; collaborative service; signal jamming; educational service; bivariate signal
History: Received: April 2015; Accepted: May 2016 by Vijay Mookerjee, after 2 revisions.

1. Introduction

Online reviews of services, in which customers rate and recommend various service providers (e.g., personal tutors, professional consultants), have become pervasive in recent years, as various ecommerce websites, user-generated-content platforms, and social networks all facilitate customers to review their service experiences (e.g., nextdoor.com, kudzu.com, elance.com). Thanks to the advancement of information systems and the ubiquitous access to the Internet, online repositories of customer reviews enable these otherwise hard-to-observe information to disseminate across spatial and temporal boundaries to become available to the general public, generating profound influences on business. Nevertheless, although a large volume of literature has been devoted to studying online reviews of products, surprisingly, online reviews of services have been left mostly untouched in the existing literature. We hence aim to fill this gap by taking an early step to theoretically explore how online reviews could impact service provision.

Within the huge service industry, one type of services that particularly draws our research interest is the educational and tutoring services (e.g., coaching students for standardized tests, teaching ice skating). As a rapidly growing industry, it is estimated that the global private tutoring market will reach $200 billion by 2020 (GIA 2014). Given the sophistication of educational activities and their lasting impacts, it is especially critical to distinguish superior providers of this type of services. A recent study shows empirical evidence that a good teacher can increase the present value of students' lifetime income by hundreds of thousands of dollars (Chetty et al. 2014). Naturally, we see a proliferation of online review platforms (e.g., nextdoor.com, myedu.com) that serve to help users identify and recommend high-quality providers of such services.

Apart from the industrial importance, the nature of the educational and tutoring services makes it especially intriguing to study the associated strategic

[Correction added on 31 May 2018, after first online publication: the affiliation and email address of Haoying Sun have been changed.]

interplay under online reviews. First, this type of services is collaborative, because the outcome of the service depends on not only the effort of the service provider but the effort of the client as well. Moreover, the service provider not only decides her own effort input but also determines the effort level of the client, at least to a certain extent (e.g., a tutor typically specifies the amount of exercise a student is required to take). As both effort levels are not directly observable to outsiders, two-dimensional moral hazard exists, and the service provider's strategic choice of these two hidden actions is thus a key issue to explore. Considering that the effort level set for the client directly affects the client's utility, will the service provider intentionally allow the client to shirk in exchange for a higher level of "happiness" and thus a better review? Consequently, will online reviews lead to strategic deterioration of the service performance? Furthermore, given the collaborative nature, the strategic correlation between the two effort decisions is another question of interest. With the client review in consideration, will the service provider invest excessive effort herself in place of the effort required from the client, or will the service provider choose to shirk as a result of the limitation on pushing the client to work diligently enough? In this sense, will the two effort choices strategically substitute or complement each other under online reviews? This study is intended to shed light on these questions.

Besides the unobservable actions, the inherent capability of a service provider to effectively deliver the service is also unobservable, making distinguishing the superior provider another important question worth examining. Different from the usual case of signaling game in which a provider can signal her superiority through observable actions, in this context, prospective clients (more generally, the market) can infer the capability of a service provider only through certain observable signals, which, in turn, can be manipulated by the hidden actions strategically chosen by the service provider. Meanwhile, the market should rationally anticipate such strategic plays and properly adjusts the belief in equilibrium. It thus constitutes the classical signal-jamming problem (Fudenberg and Tirole 1986). Being informative about the inherent capability of the provider, online reviews can certainly serve as such a signal, whose effects on service provision hence deserve careful analysis. Another natural candidate for such a signal is the outcome of the service itself.[1] Depending on their availability, either the service outcome or the client review, or both, can be observed by the market. We thus use the service outcome as a benchmark case to contrast with the cases when the market observes the client review either alone or along with the service outcome. We are interested in comparing the resulting performance of the strategic service provision and the distinguishability of the service provider's underlying capability.

While motivated from the context of educational and tutoring services, our model and results have general implications for other types of services with similar features, as long as the service is collaborative in nature, the service provider can determine (at least to a certain extent) the effort level of the client, and the service outcome is objective. Examples include consulting services (e.g., IT consulting projects requiring collaboration from both parties), professional services (e.g., taxation and auditing, legal services), and health care services (e.g., rehabilitation and physical therapy).

In this study, we develop a signal-jamming model and compare the equilibrium effort levels when the service outcome is observed as a signal and when the client review serves as a signal. We find that the review generally leads to a lower effort level for the client, as the provider is concerned about the client's disutility of exerting effort to be negatively reflected in the review. Meanwhile, the review does lead to an efficient effort level for the client, that is, the same effort level the client would have chosen for himself. For the provider's own effort decision, we identify two effects determining the provider's incentive for effort investment, namely, the *affirmative effect* and the *informative effect*. The former refers to how likely the effort input can result in a positive signal, whereas the latter means how convincingly a positive signal implies the superiority of the provider. We show that under online reviews, two factors—the additional uncertainty of the review and the lower effort from the client—weaken both affirmative and informative effects, resulting in a lower effort level of the provider. In this sense, the client review results in a sacrifice of the service outcome in favor of the client's utility, which leads to a lower degree of distinguishability of the provider type in equilibrium.

We further extend the analysis by allowing the client to observe private information about the provider type upon service completion and to incorporate such information into his review. As a result, the informative effect is strengthened, and the provider is better motivated to work diligently. Interestingly, when the review is highly informative, the provider's two-dimensional effort decisions may be strategic substitutes in that the provider requires less effort from the client compared to the benchmark case, yet she invests excessive effort of her own. As we show, with a reasonable level of informativeness, online reviews could lead to favorable equilibrium performance in multiple dimensions: higher service effectiveness, more efficient effort level for the client, and better distinguishability of the provider type.

We also investigate the case in which the client review is observable along with the service outcome. Counterintuitively, we demonstrate that richer information may not necessarily generate favorable results as an outcome of strategic interplay. When both the review and the outcome are observable to the market, the service provider surprisingly lacks sufficient incentive to devote effort, simply because the bivariate signal structure makes a successful service outcome no longer rewarding if the review is negative. As a result, both signals being available at the same time could lead to inferior distinguishability of provider types.

To the best of our knowledge, this study is the first theoretical study connecting online reviews with collaborative services. From modeling perspective, a distinctive aspect of our study is that we introduce the signal-jamming framework into the context of online reviews and service management, which enables us to explicitly model the review generating and information updating processes and to endogenously incentivize effort input without external agency contracts. Our results also enrich the information economics by illustrating how strategic interplay under a complex information structure can lead to surprising economic consequences. On the substantive front, the findings in this study provide rich implications for various stakeholders, including information systems designers, top management and authorities, service providers, and clients who write reviews, as we discuss in detail in the end of the article.

2. Literature Review

As we study the unique information structure associated with online reviews in the context of collaborative services, our modeling approach is rooted in the classical signal-jamming framework (Fudenberg and Tirole 1986, Holmstrom 1999). The problem we study and the modeling approach we adopt are thus distinct from what is typically found in the existing literature on both online reviews and collaborative services, as we briefly review below.

The existing literature on online reviews predominantly focuses on customer reviews of products. There is a large volume of literature empirically investigating how different aspects of online reviews affect the sales of various products such as books (e.g., Chevalier and Mayzlin 2006), movies (e.g., Liu 2006), video games (e.g., Zhu and Zhang 2010), and digital cameras (e.g., Gu et al. 2012). Meanwhile, theoretical studies develop analytical models to examine how product reviews could affect product sales (e.g., Li and Hitt 2010), firm strategies (e.g., Chen and Xie 2008), and market competition (e.g., Kwark et al. 2014). Recent literature (e.g., Jiang and Guo 2015) further explores the optimal design of review systems with endogenous consideration of product features such as product mainstream level and consumer misfit cost. Studying reviews of services, we thus differ from the extant literature in the research context and the main focus.

In modeling the role of online reviews, the aforementioned theoretical papers typically view it as conveying product information and generally take one of the two approaches. One class of approach (e.g., Chen and Xie 2005, 2008, Kwark et al. 2014) considers reviews as an exogenous information source providing additionally accurate information about product quality and/or fit, where the review generating process is left exogenous. Another class of approach (e.g., Hao et al. 2011, Jiang and Guo 2015, Sun 2012) models the generation of reviews as a deterministic function of customer utility involving observable information (e.g., price) and individual reviewer heterogeneity. As a result, rational prospective customers can back out the product information (e.g., quality, fit) based on the observable information (e.g., valence and variance of reviews, price). Differently, we not only endogenously model the review generating and information updating processes, but also account for the effects of hidden strategic actions and possible noises. In consequence, how accurately the review can convey information about the latent provider ability is a strategic outcome endogenously dependent on the information structure, and reviews may not necessarily lead to more accurate information in equilibrium. In this sense, our modeling approach is close to Kuksov and Xie (2010). They consider a two-period model of product reviews conveying information about the latent product quality, which both the firm and customers are unsure of. The firm strategically decides the price and the "frill", both unobservable to the second-period customers, to influence the reviews given by the first-period customers. Besides the different research context, the information structure studied in our study is also different from theirs in that we introduce an additional layer of uncertainty into the review, so it depends on another random variable (i.e., the service outcome). Consequently, we can further study the effects when both signals are observable and discover interesting equilibrium results.

More broadly, our study is also related to another stream of literature on reputation (e.g., Bakos and Dellarocas 2011, Kreps and Wilson 1982, Milgrom and Roberts 1982), which takes the view that reputation is built and maintained as an indication of credible commitment (i.e., being locked into playing a certain strategy) in repeated games. In contrast, similar to the aforementioned theoretical papers on product reviews, our study emphasizes the role of reviews as revealing additional information about underlying key characteristics (i.e., ability). In addition to the different fundamental views, other distinctions include: we endogenously model the generation of reviews, which depend on service outcomes and strategic actions as well, whereas reputation is typically considered as a direct collection or an externally-specified transformation of past performances; we compare equilibrium results under different signals, which, in a broad sense, provides richer implications for the effects of different reputation mechanisms; moreover, we link online reviews to collaborative services, filling a gap between these two important areas.

Starting from Karmarkar and Pitbladdo (1995), collaborative (or coproductive) services have been attracting increasing research interest in the literature of operations management (e.g., Plambeck and Taylor 2006, Roels et al. 2010, Spohrer and Maglio 2008, Xue and Field 2008) as well as information technology outsourcing (e.g., Bhattacharya et al. 2014, Demirezen et al. 2013). The extant literature in this domain predominantly follows the principal-agent contract-design framework. In contrast, as we study a new problem in this context, we build upon the theories of career concerns (Holmstrom 1999) and allow the service provider to be incentivized endogenously without externally stipulated agency contracts or in the face of possible incomplete contracts. Another difference is that the service provider also determines the client's effort in addition to her own in our context, which increases the dimension of the provider's strategy space and leads to interesting strategic correlation.

3. The Baseline Model

3.1. Model Setup

We consider one service provider and one client in the market.[2] The service is collaborative, and the outcome of the service depends on the effort levels of both the service provider and the client. We denote the effort level of the service provider as x and that of the client as y, where both are bounded and normalized such that $x, y \in [0, 1]$. In addition, the service outcome is also influenced by the innate ability of the service provider to deliver the service effectively, which we refer to as the *type* of the provider. We assume there are two types of the service provider, high (H) type or low (L) type, where a high-type provider is more likely to deliver the service effectively than a low-type provider given the same effort levels. Besides these focal factors, many other factors not directly pertinent to our interest may also affect the outcome. Therefore, the outcome is probabilistic by its nature. We hence model the outcome of the service, V, as a random variable taking a binary value, 0 or 1, where 1 represents success (e.g., passing a standardized test, achieving the preset goal) and 0 represents failure. The value of V depends on a latent variable v, which represents the combination of all possible factors affecting the service outcome. If $v > 0$, we have $V = 1$; if $v \leq 0$, we have $V = 0$. Specifically, v is modeled as follows.

$$v(\xi; \theta, x, y) = \mu_\theta(x + y) + \xi, \quad \theta \in \{H, L\}. \qquad (1)$$

Here, μ_θ reflects the innate ability of the θ-type provider, and $0 < \mu_L < \mu_H$. Without loss of generality, we normalize $\mu_H = 1$ and let $\mu_L = \mu \in (0, 1)$. The term $(x + y)$ represents the combined effort from both parties, which captures the collaborative nature of the service and is commonly used in the literature (e.g., Karmarkar and Pitbladdo 1995). The last term in (1), ξ, represents other idiosyncratic factors that could possibly affect the service outcome. It is a random variable with the cumulative distribution function (cdf) $G(\xi)$. To derive clean analytical results, we let ξ follow a uniform distribution over the support $[-2, 0]$, so $G(\xi) = \frac{\xi+2}{2}$. The support is chosen such that when the deterministic term $\mu_\theta(x + y)$ takes its minimum value 0, we have $v(\xi) \leq 0$ for any possible ξ; when $\mu_\theta(x + y)$ takes its maximum value 2, we have $v(\xi) \geq 0$ for all ξ's. Notice that neither the distribution nor the support of ξ is critical for our results to hold. As we show in Appendix S2, under different distributions (e.g., beta distribution, normal distribution) or expanded supports (either finite or infinite), our main results remain robust.

As a result, the probability of the outcome being a success, given the true type of the service provider and the effort levels of both parties, can be derived as

$$\begin{aligned}\Pr\{V = 1|\theta, x, y\} &= \Pr\{v(\xi) > 0|\theta, x, y\} \\ &= 1 - G(-\mu_\theta(x + y)) \\ &= \frac{1}{2}\mu_\theta(x + y), \quad \theta \in \{H, L\}.\end{aligned} \qquad (2)$$

After exerting the effort y and observing the realized outcome V, the client derives net utility $U(V, y) = u\mathbb{1}_{V=1} - wy^2$, where $\mathbb{1}_{V=1}$ is an indicator function such that $\mathbb{1}_{V=1} = 1$ if $V = 1$, and $\mathbb{1}_{V=1} = 0$ otherwise. In other words, the client derives positive utility gain u if the outcome turns out a success and zero utility gain if a failure; meanwhile, the utility cost from exerting effort is wy^2. Here, w is the client's effort cost coefficient ($0 < w < u$), and the convex functional form captures the increasing marginal disutility as the effort level increases. Without loss of generality, we normalize $u = 1$, and thus $w \in (0, 1)$.

Given the realized utility, the client then writes an online review of the service provider. Considering that various latent factors could all affect the review (e.g., leniency, preference), we adopt a similar structure as the outcome and model the review, R, as a random variable taking a binary value: $R = 1$ means the client recommends the service provider, and $R = 0$ means he does not recommend. The binary structure reflects the commonly observed rating mechanism used in online reviews and is also common in the literature (e.g., Kuksov and Xie 2010). For easy reference, we refer to $R = 1$ as a positive review and $R = 0$ a negative review. The value of R depends on a latent variable ϱ, which can be interpreted as the client's overall recommendation intention. If $\varrho > 0$, we have $R = 1$; if $\varrho \leq 0$, we have $R = 0$. As theories and evidences on rater behavior suggest, reviews depend on the net utility of the reviewers along with other idiosyncratic factors.[3] We hence model ϱ as follows.

$$\varrho(\varepsilon; V, y) = U(V, y) + \varepsilon, \tag{3}$$

where ε stands for possible idiosyncratic factors that could affect the review. It is a random variable with cdf $F(\varepsilon)$. Similarly, for neat analytical results, we let ε follow a uniform distribution over the support $[-1, w]$, so $F(\varepsilon) = \frac{\varepsilon+1}{w+1}$. The support is chosen under the same rationale as the case of ξ: when the client's utility $U(V, y)$ takes its minimum value $-w$ (i.e., when $V = 0$ and $y = 1$), $\varrho(\varepsilon) \leq 0$ for any possible ε; when $U(V, y)$ takes its maximum value 1 (i.e., when $V = 1$ and $y = 0$), $\varrho(\varepsilon) \geq 0$ for all ε's. Again, neither the distribution nor the support of ε is critical. Our main results remain robust when extended to different distributions or expanded supports, as we show in detail in Appendix S2.

As a result, the probability of the client writing a positive review, given the outcome V and the client's effort y, can be derived as

$$\begin{aligned}\Pr\{R = 1|V, y\} &= \Pr\{\varrho(\varepsilon) > 0|V, y\}\\ &= 1 - F(-U(V, y))\\ &= \frac{\mathbb{1}_{V=1} - wy^2 + w}{1 + w}.\end{aligned} \tag{4}$$

In determining the payoff of the service provider, we model it endogenously in a general approach. We let the market form a rational belief about the true type of the service provider and reward the provider accordingly. After observing the available signal(s) (e.g., the outcome V, the review R), the market forms a posterior belief α, which equals the posterior probability of the service provider being a high type given the observed signal(s). We will articulate the formation of α in detail in subsection 3.2. Given the belief α, the market rewards the provider with a payoff equal to $\alpha m_H + (1 - \alpha)m_L$, where $m_L < m_H$ so it is more rewarding to be viewed as a high-type provider. Notice that the payoff from the market is thus an abstract form of the present value of the provider's future payoff. It can be interpreted broadly, for example, as the monetary benefit rationally offered by a third-party authority, or as the willingness-to-pay of future clients. Additionally, we let the service provider's utility cost of exerting effort take the similar convex form as that of the client, that is, cx^2, where c is the provider's effort cost coefficient ($0 < c < m_H$). Once again, to simplify notation, we normalize $m_H = 1$ and $m_L = 0$ without loss of generality. As a result, the service provider's net payoff, given the market belief α, can be written as

$$\pi(x; \alpha) = \alpha m_H + (1 - \alpha)m_L - cx^2 = \alpha - cx^2. \tag{5}$$

3.2. Equilibrium Concept

As this study focuses on the type of collaborative services that the service provider not only determines her own effort but also the effort level of her client, we consider that the provider decides *both* effort levels x and y. (In subsection 6.1, we extend the model to the general case that both the provider and the client jointly determine the client's effort level y, and each has a varying degree of influence. As we show, the main results continue to hold.) As a result, the actual strategic players in the model are the service provider and the market. From the market's standpoint, the effort levels x and y chosen by the service provider are not observable.[4] Instead, the market can infer the true type of the provider only through observable information, such as the outcome V, the review R, or both, which, in turn, depends on the hidden actions taken by the provider. Such an information structure constitutes a typical signal-jamming game (Fudenberg and Tirole 1986). Following this framework, we let the market and the provider share common knowledge about all model parameters, including the prior belief about the true provider type. In other words, the true type of the provider, which is essentially her underlying capability to deliver the service effectively, is not directly observable to either the market or the provider herself, whereas they share a common prior (probabilistic) belief about it. It is common in the classical literature on job performance and career concerns to assume that a player cannot observe her latent characteristic and share the same prior belief with the market (e.g., Holmstrom 1999). For simplicity, we use the uninformative prior such that the prior probability of the service provider being a high type is $\frac{1}{2}$, that is, $\Pr\{\theta = H\} = \frac{1}{2}$. Nevertheless, choosing a different prior other than $\frac{1}{2}$ will not qualitatively change our main results.

The equilibrium concept we use is *perfect Bayesian equilibrium*. In equilibrium, on one hand, the service provider's choice of effort levels are optimal given the market's belief; on the other hand, the market's belief is rational (according to Bayes' rule) given the provider's effort choice. In the subsequent analysis, we will study separate scenarios in which the market can observe the outcome V only, the review R only, or both the outcome and the review. We use S to represent the signal that the market can observe. Depending on different scenarios, S can be V, R, or $\{V,R\}$. Note that S is thus a random variable (or a pair of random variables) taking binary values.

Upon observing the realized value of S, the market rationally forms posterior belief about the true type of the service provider. Specifically, the market's belief α_S^* equals the posterior probability of the service provider being a high type, conditional on the observed value of signal S and given the provider's equilibrium effort levels x^* and y^*. By Bayes' rule,

$$\begin{aligned}\alpha_S^* &= \Pr\{\theta = H|S, x^*, y^*\}\\ &= \Pr\{S|\theta = H, x^*, y^*\} \cdot \Pr\{\theta = H\}/\\ &\quad \Big(\Pr\{S|\theta = H, x^*, y^*\} \cdot \Pr\{\theta = H\}\\ &\quad + \Pr\{S|\theta = L, x^*, y^*\} \cdot \Pr\{\theta = L\}\Big).\end{aligned} \tag{6}$$

Meanwhile, anticipating the market belief α_S^*, the service provider makes an upfront decision of the optimal effort levels such that x^* and y^* maximize the *ex ante* expectation of her net payoff (i.e., taking expectation over the possible realization of the signal S conditional on the effort levels). Specifically,

$$\begin{aligned}(x^*, y^*) &= \arg\max_{0 \le x,y \le 1} E_\theta\{E_S[\pi(x; \alpha_S^*)|\theta, x, y]\}\\ &= \arg\max_{0 \le x,y \le 1} E_\theta\{E_S[\alpha_S^*|\theta, x, y]\} - cx^2,\end{aligned} \tag{7}$$

where the second equality holds by substituting in (5). Notice that although the effort level of the client y does not directly enter the provider's payoff function (5), it does affect the provider's expected payoff by influencing the realization of the signal S. Therefore, the service provider needs to optimize both x and y altogether, which introduces the strategic interaction between these two decisions.

Altogether, a perfect Bayesian equilibrium of this signal-jamming game is hence a fixed point $\{x^*, y^*, \alpha_S^*\}$ that solves (6) and (7) simultaneously.

4. Analysis and Results

In this section, we first solve the equilibrium under separate scenarios in which (i) the market observes the outcome V alone as a signal, (ii) the market observes the review R alone as a signal, or (iii) the market observes both the outcome and the review. We then derive formal results comparing the equilibrium effort levels under these different scenarios, which further allow us to examine the equilibrium outcomes in various dimensions such as the effectiveness of service, the efficiency of the effort choice, and the distinguishability of the true type of the service provider.

4.1. Outcome as a Signal

We first examine the benchmark scenario in which the market can observe the outcome V, whereas the review R is unavailable, so $S = V$. To differentiate from the other two scenarios, we use the superscript V to indicate the equilibrium solutions in this scenario instead of asterisk.

In equilibrium, anticipating the equilibrium effort levels chosen by the service provider, x^V and y^V, the market rationally updates its posterior belief according to (6), which can be derived as $\alpha_1^V = \Pr\{\theta = H|V = 1, x^V, y^V\} = \frac{1}{1+\mu}$ and $\alpha_0^V = \Pr\{\theta = H|V = 0, x^V, y^V\} = \frac{2-(x^V+y^V)}{4-(1+\mu)(x^V+y^V)}$. Given the market belief, the service provider optimizes the effort levels to maximize her expected payoff, $E\pi$, according to (7), where $E\pi = \alpha_0^V + \frac{1+\mu}{4}(\alpha_1^V - \alpha_0^V) \cdot (x + y) - cx^2$. Notice that the market forms its belief without observing the actual effort, so the equilibrium concept requires that α_1^V and α_0^V, and hence all x^V's and y^V's included, be treated as constants when the provider optimizes x and y. Denote $\Delta\alpha^V(x^V, y^V) \equiv \alpha_1^V - \alpha_0^V$, we can write the first derivative with respect to x and y as

$$\frac{\partial E\pi}{\partial x} = \frac{1+\mu}{4}\Delta\alpha^V(x^V, y^V) - 2cx, \tag{8}$$

$$\frac{\partial E\pi}{\partial y} = \frac{1+\mu}{4}\Delta\alpha^V(x^V, y^V). \tag{9}$$

Notice that $\frac{\partial E\pi}{\partial y}$ is a constant, and it is easy to show that this constant is always positive, that is, $\Delta\alpha^V(x^V, y^V) > 0$ for $\forall x^V, y^V$ (see Lemma A.1 in Appendix S1 for the proof). Therefore, the provider's expected payoff strictly increases in y. As a result, the optimal choice of y is the upper bound of the client's effort range, that is, $y^V = 1$. Consequently, the optimal choice of the service provider's own effort level, x^V, is the fixed-point solution to the first order condition $\frac{1+\mu}{4}\Delta\alpha^V(x^V, 1) - 2cx^V = 0$ by (8), whenever such an interior solution exists within [0, 1].

Following the procedure outlined above, we can formally derive the equilibrium effort levels when the outcome is observed as a signal, as follows.

Proposition 1. *When the market observes the outcome V alone as a signal, for $c > \frac{1}{8}$, there is a unique equilibrium in which the service provider's effort* $x^V = \frac{3-\mu}{2(1+\mu)} - \sqrt{\frac{(3-\mu)^2}{4(1+\mu)^2} - \frac{1-\mu}{4c(1+\mu)}}$ *and the client's effort* $y^V = 1$.

PROOF. All proofs are detailed in Appendix S1. □

Proposition 1 shows that when only the outcome is used as a signal, in equilibrium, the service provider sets the highest possible effort level for the client. Note that $\alpha_1^V > \alpha_0^V$, which means a successful outcome gives the market more confidence to believe the service provider is a high type and is thus more rewarding to the provider. Hence, when the review is unavailable and the provider does not need to take account of the client's utility, she pushes the client's effort to the maximum in order to maximize the probability of getting a successful outcome. For the provider's own effort, she trades off between the expected gain from market recognition and the utility cost of exerting effort. As a result, the optimal effort level of her own, x^V, is an interior solution between 0 and 1.

It is worth emphasizing that we prove in Proposition 1 that there exists a *unique* equilibrium for the proposed signal-jamming game generically (i.e., as long as $c > \frac{1}{8}$). In the extreme case when the service provider's effort cost coefficient is very small (i.e., $c < \frac{1}{8}$), multiple equilibria could arise, among which $x^V = y^V = 1$ is always an equilibrium. We find similar results for the other scenarios as well. The multiple equilibria yield no closed-form full characterization and involve discussion on equilibrium refinement, whereas the main results remain the same, as we can show. Therefore, to deliver neat results with clear implications and to avoid technical distraction, in the main body of the study, we focus on the generic cases when the service provider's effort cost coefficient is not too small. In Appendix S3, we characterize the equilibria for small c's in detail and show the main results extend robustly.

COROLLARY 1. *The equilibrium effort level of the service provider, x^V, is decreasing in her effort cost coefficient c, and increasing in the difference between the two types of providers $1 - \mu$ (i.e., decreasing in μ).*

Corollary 1 shows some interesting comparative statics of the equilibrium effort. For example, as μ decreases and hence the capability gap between the two types of providers enlarges, the service provider is willing to work more diligently. This is because as the gap between the two types of providers increases, the outcome becomes more revealing, so a successful (or failed) outcome signals more convincingly to the market that the service provider is a high (or low) type. As a result, the marginal return for the service provider to devote more effort to increase the probability of a successful outcome also goes up, which leads to a higher equilibrium effort level x^V.

4.2. Review as a Signal

We next analyze the scenario in which the review R is available, whereas the outcome V, although still observable to the service provider and the client, is unobservable to the market. Thus, $S = R$. Again, to differentiate from the other scenarios, we use the superscript R to indicate the equilibrium solutions in this scenario instead of asterisk.

The equilibrium can be derived following a similar approach as described in subsection 4.1. A main difference lies in formulating the likelihood function, $\Pr\{R|\theta, x^R, y^R\}$. Taking the expectation of the (conditional) probability of receiving a positive review (given the outcome V) in (4) over V, we have

$$\Pr\{R|\theta,x^R,y^R\} = \Pr\{R|V=1,y^R\}\Pr\{V=1|\theta,x^R,y^R\} + \Pr\{R|V=0,y^R\}\Pr\{V=0|\theta,x^R,y^R\}. \quad (10)$$

Substituting (2) and (4) into (10), we can derive the likelihood function, which leads to the equilibrium market belief, $\alpha_1^R(x^R, y^R)$ and $\alpha_0^R(x^R, y^R)$, according to the Bayes' rule in (6). Given the market belief, similar as in subsection 4.1, the service provider optimizes both x and y, and the equilibrium effort levels are the fixed-point solution. Denoting $\Delta\alpha^R(x^R, y^R) = \alpha_1^R - \alpha_0^R$, we can summarize the equilibrium effort levels when the review is observed as a signal as follows.

PROPOSITION 2. *When the market observes the review R alone as a signal, for $c > \frac{1}{2}M(1)$, there is a unique equilibrium in which the client's effort $y^R = \min\{\frac{1+\mu}{8w}, 1\}$ and the service provider's effort x^R is the unique solution to $M(x) = 2cx$ within the range of $x \in [0, 1]$, where $M(x) \equiv \frac{1+\mu}{4(1+w)}\Delta\alpha^R(x, y^R)$.*

Note that the function $M(x)$ defined in Proposition 2 originates from the first-order condition with respect to x such that $M(x^R) = 2cx^R$. As detailed in Lemma A.2 in Appendix S1, the monotonicity of $M(x)$ and $\frac{\partial^2}{\partial x^2}M(x)$ results in a single crossing of $M(x)$ and $2cx$ for $x \in [0, 1]$ when c is not too small, which ensures the existence and uniqueness of the equilibrium effort x^R.

As Proposition 2 implies, when the review is observed as the signal, the service provider no longer always chooses the maximum effort level for the client. As long as the client's cost coefficient w is not too small (i.e., $w > \frac{1+\mu}{8}$), the optimal effort of the client y^R is set as an interior solution less than 1.

COROLLARY 2. *The equilibrium effort of the service provider, x^R, is decreasing in her effort cost coefficient c and the client's effort cost coefficient w; the equilibrium effort of the client, y^R, is (weakly) decreasing in w and* *independent of c. While y^R is (weakly) increasing in μ, x^R is decreasing in μ when $w \leq \frac{1+\mu}{8}$ and may change non-monotonically in μ when $w > \frac{1+\mu}{8}$.*

Corollary 2 shows the comparative statics of the equilibrium effort. An interesting result is that as the client's effort becomes more costly, not only the client's equilibrium effort y^R decreases, the provider's equilibrium effort x^R also decreases. It thus implies some degree of complementarity between the equilibrium efforts of the provider and the client, as we will further explore. The possible non-monotonicity of x^R in μ is also worth discussing. When $w \leq \frac{1+\mu}{8}$ (so $y^R = 1$), x^R is decreasing in μ, similar to Corollary 1. When $w > \frac{1+\mu}{8}$ (so $y^R = \frac{1+\mu}{8w} < 1$), on one hand, a higher μ reduces the difference between the two provider types and thus the informational value of a positive signal, which lowers the provider's effort incentive; on the other hand, the provider can require a higher effort level from the client as μ increases, which, in turn, strengthens the provider's own effort incentive. The two counteracting factors hence cause the possible non-monotonicity, and x^R may first increase and then decrease in μ as μ goes up.

4.3. Outcome and Review as Signals

When the market can observe both the outcome V and the review R, both pieces of information will be incorporated into the rational expectation about the true type of the service provider. Under the baseline model, because the review does not convey additional information regarding the true type of the provider beyond what is already conveyed by the outcome, the equilibrium when both V and R are observable turns out to be the same as when only V is available as the signal (i.e., the equilibrium in subsection 4.1).

Formally, such an equivalence originates from the conditional independence of R on θ given V, that is, $\Pr\{R|V, \theta, x, y\} = \Pr\{R|V, y\}$. In other words, given the realization of the outcome V, the probability of receiving a positive review no longer depends on the true type of the provider. As a result, $\Pr\{V, R|\theta, x, y\} = \Pr\{R|V, y\} \cdot \Pr\{V|\theta, x, y\}$, and in Bayesian updating according to (6), the term $\Pr\{R|V, y\}$ cancels out, and we have $\Pr\{\theta = H|V, R, x, y\} = \Pr\{\theta = H|V, x, y\}$. It means the posterior market belief when both V and R are observed will be the same as when only V is observed, independent of the value of R. As a result, the equilibrium effort will be exactly the same as the x^V and y^V derived in Proposition 1.

In Section 5, we will extend the baseline model to incorporate additional information regarding the true type of the provider into the review. In that case, observing both signals will lead to different equilibrium effort, which can deliver richer implications regarding how client reviews could affect collaborative service provision.

4.4. Equilibrium Comparison

Having obtained the equilibrium in the two scenarios where either the outcome or the review is observed as a signal, we next examine how the equilibrium effort differs across the two scenarios. We also compare the equilibrium performance of the service in terms of effectiveness, efficiency, and distinguishability.

PROPOSITION 3. *The equilibrium effort levels of both the service provider and the client are lower when the review is observed as a signal than when the outcome is observed as a signal. Specifically, for $c > \frac{1}{8}$, $x^R < x^V$, and $y^R \leq y^V$ with strict inequality when $w > \frac{1+\mu}{8}$.*

Proposition 3 shows that with the review observed as a signal, the service provider generally sets a lower effort level for the client and devotes less effort herself as well, in comparison to the case when the outcome is observed as a signal. To understand the result that $y^R \leq y^V$, recall that the client gives the review based on his net utility, which is increased by a successful outcome and decreased by his disutility from exerting effort. Therefore, to maximize the probability of receiving a positive review, the provider faces the trade-off between a higher success probability and elevated effort disutility of the client. The result in Proposition 3 shows that the disutility concern generally dominates the success motivation, so the provider typically requires lower effort from the client compared to the case when only the outcome matters. Therefore, $y^R < y^V$ unless the client's effort is most costless (i.e., w is small).

The comparison regarding the equilibrium effort of the provider's own is more subtle and intriguing to explain. When optimizing her own effort, the service provider faces the trade-off between convincing the market of her capability and exerting costly effort. In determining the marginal benefit of improving market recognition from effort investment, two effects come into play: The first is the *affirmative effect*, that is, how greatly an increased level of effort could improve the chance of attaining a positive signal. The second is the *informative effect*, that is, how much information a positive signal conveys to the market regarding the provider being a high type. When the review serves as a signal, there are two factors that reduce both the affirmative and the informative effects. First, compared to the outcome, the review is influenced by additional factors such as the client's effort cost and many other unobservables, all of which bring more uncertainty and noise into the realization of the review value. Second, the equilibrium effort level of the client is lower when the review serves as a signal, as just discussed.

Less effort from the client reduces the difference between the two types of providers in terms of the probability of attaining a positive signal. Driven by these two factors (i.e., additional uncertainty and less effort from the client), when the review serves as a signal, the provider's effort is less rewarding in generating a positive signal, and moreover, a positive signal becomes less informative about the true type of the provider. As a result, both the affirmative and the informative effects are reduced, so the marginal benefit of the provider's effort investment decreases, leading to a lower level of x^R in equilibrium.

It is worth noting that in the baseline model, both the affirmative and the informative effects change in the same direction across the two scenarios, causing a decisive comparison result. In addition, compared to the outcome, the review leads to lower effort levels of both the provider and the client. In this sense, the effort choices exhibit a strategic-complement pattern in equilibrium. In Section 5, when we introduce additional information about the provider type into the review, interestingly, the review may change the affirmative and informative effects in the opposite directions, and the provider's and the client's equilibrium effort levels may be strategic substitutes instead.

Now that the review leads to less equilibrium effort of both the provider and the client, we can immediately compare the effectiveness of service, which we measure using the expected outcome in equilibrium, defined as

$$EV(x^*, y^*) = E_\theta\{E[V|\theta, x^*, y^*]\}. \tag{11}$$

Proposition 4. *The effectiveness of service is lower when the review is observed as a signal than when the outcome is observed as a signal in that $EV(x^R, y^R) < EV(x^V, y^V)$.*

We are also interested in investigating the efficiency of the equilibrium effort the provider sets for the client in different scenarios.[5] We define the *efficient* effort level of the client, y^e, as the one that maximizes the client's expected net utility, that is, the optimal effort level from the client's own perspective. Specifically,

$$\begin{aligned} y^e &= \arg\max_{0\le y\le 1} E_{\theta,V}[U(V,y)] \\ &= \arg\max_{0\le y\le 1} EV(x,y) - wy^2, \end{aligned} \tag{12}$$

where EV is defined in (11) (Recall that the client's utility gain from a successful outcome, u, is normalized to 1 for simplicity). It is easy to solve that $y^e = \min\{\frac{1+\mu}{8w}, 1\}$, which is independent of the provider's effort x. Interestingly, the equilibrium effort level set for the client when the review is observed as a signal, y^R, turns out to be efficient.

Proposition 5. *The client's effort is at the efficient level when the review is observed as a signal, whereas the client's effort generally exceeds the efficient level when the outcome is observed as a signal, that is, $y^R = y^e \le y^V$ with strict inequality when $w > \frac{1+\mu}{8}$.*

Combining the results of Propositions 4 and 5, it becomes clear that when the review serves as a signal, the equilibrium output is less success oriented; instead, it caters to the client to optimize his effort input by sacrificing the effectiveness to a certain degree. In contrast, when the outcome serves as a signal, the client is pushed to work beyond his efficient effort level.

Another dimension we are interested in exploring is the distinguishability of provider types. We want to examine how accurately the market can tell the true type of the provider in equilibrium under different signals. The natural means to measure the accuracy of the market's posterior belief is the usual type I and type II errors. In our context, we define the equilibrium type I error, ρ_1^*, as the *ex ante* probability of mistakenly believing a high-type provider as a low type, given the equilibrium effort levels. The equilibrium type II error, ρ_2^*, is thus the *ex ante* probability of mistakenly believing a low-type provider as a high type. By *ex ante*, we mean taking the expectation over possible realized values of the signal S. We can formulate the two types of errors as follows.

$$\rho_1^* = E_S\left[1 - \alpha_S^*|\theta = H, x^*, y^*\right] \tag{13}$$

$$\rho_2^* = E_S\left[\alpha_S^*|\theta = L, x^*, y^*\right] \tag{14}$$

It is noteworthy that because the prior belief is symmetric and the posterior belief is formed rationally based on the Bayes' rule, it can be shown that the equilibrium type I and type II errors happen to be the same in our context. This coincidence leads to a simple and neat measure of the equilibrium distinguishability of provider types, which we refer to as the equilibrium judgment error ρ^*, where $\rho^* = \rho_1^* = \rho_2^*$.[6]

Figure 1 illustrates the equilibrium judgment errors under different signals with different parameter values. As we can see, $\rho^R > \rho^V$ in general, which means the review leads to a higher equilibrium judgment error than the outcome. This result can be explained by the aforementioned two factors that reduce the informative effect: the review brings in additional uncertainty and also causes a lower equilibrium effort level of the client. It can be clearly seen from Figure 1 how less effort from the client drives a higher equilibrium judgment error: as the client's effort cost coefficient w goes up, the client's effort y^R decreases, and ρ^R increases in consequence. Another noteworthy pattern in Figure 1 is the increase of judgment error in μ.

Figure 1 Illustration of Equilibrium Judgment Error (c = 0.15)

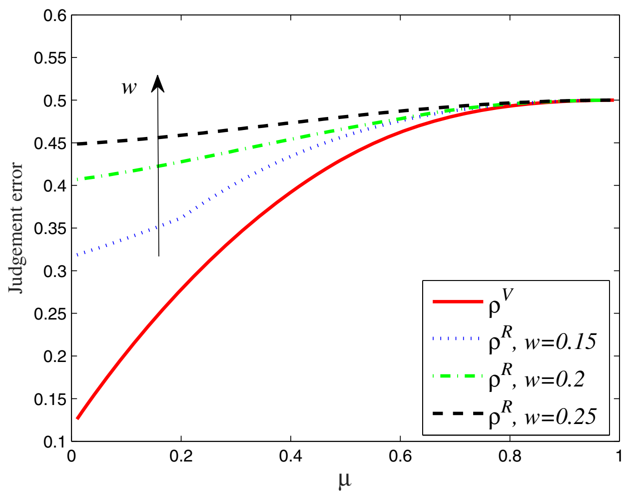

As μ increases and gets closer to 1, the low-type provider is closer to the high-type provider in terms of the innate ability of successfully delivering the service. As a result, it becomes more difficult to tell apart the two types of providers, leading to a higher judgment error in equilibrium.

5. Review with Private Information

In the baseline model, the review R does not contain any additional information about the provider type beyond what has already been included in the outcome V. As discussed in subsection 4.3, this property leads to unchanged equilibrium results when the review is observed along with the outcome. In this section, we extend the baseline model by letting the review contain additional information beyond the outcome. The purpose of this extension is twofold. First, we want to show to what extent the main results from the baseline model are robust and if any new results with richer implications will arise. Second, we are interested in shedding light on how the review can further impact the collaborative service provision when the outcome is already observable as a signal.

Following the same setup of the baseline model, now we assume that upon the completion of the service, the client can observe a private signal, $\tilde{\theta}$, that suggests the true type of the service provider. This signal is noisy in that it may correctly reflect the true provider type only with certain probability. Let $\Pr\{\tilde{\theta} = H | \theta = H\} = \sigma_H$, and $\Pr\{\tilde{\theta} = H | \theta = L\} = \sigma_L$, $\sigma_H, \sigma_L \in (0, 1]$. In other words, the client may perceive a high-type (or low-type) provider as a high type with probability σ_H (or σ_L). We consider this noisy signal (weakly) informative so that $\sigma_H \geq \sigma_L$, that is, a high-type provider is more likely to be perceived by the client as a high type than a low-type provider. It is then reasonable to assume that the client will recommend the service provider only if he perceives her as a high type. Thus, we can modify the probability of the client giving a positive review, previously defined in (4), as

$$\Pr\{R = 1 | \theta, V, y\} = \sigma_\theta \frac{\mathbb{1}_{V=1} - wy^2 + w}{1 + w}. \tag{15}$$

Equation (15) implies that if the client receives a positive signal suggesting the provider is a high type, he then gives a positive review with a probability that is increasing in his net utility as modeled previously. In this sense, the baseline model is a special case with $\sigma_H = \sigma_L = 1$.

Apparently, the change in the model does not affect the equilibrium results when only the outcome is observed by the market. Therefore, the results in subsection 4.1 continue to hold under this extended model. On the other hand, it may affect the equilibrium results when the review alone is available as the signal and when both the review and the outcome are observable to the market. In what follows, we present the results in these two scenarios under this extended model. Throughout this section, we add a tilde ($\sim$) to indicate the equilibrium solutions under this extended model.

5.1. Review as a Signal

The equilibrium efforts can be derived following the same approach outlined in subsection 4.2. Denoting $\Delta\tilde{\alpha}^R(\tilde{x}^R, \tilde{y}^R) = \tilde{\alpha}_1^R - \tilde{\alpha}_0^R$, we are able to solve the equilibrium under this extended model as follows.

PROPOSITION 6. *When the market observes the review R alone as a signal, for any* $0 < \sigma_L \leq \sigma_H \leq 1$, *if* $c > \frac{1}{2}\tilde{M}(1)$, *there is a unique equilibrium in which the client's effort* $\tilde{y}^R = \min\{\frac{\sigma_H + \sigma_L \mu}{4w(\sigma_H + \sigma_L)}, 1\}$ *and the service provider's effort* $\tilde{x}^R$ *is the unique solution to* $\tilde{M}(x) = 2cx$ *within the range of* $x \in [0, 1]$, *where* $\tilde{M}(x) \equiv \frac{\sigma_H + \sigma_L \mu}{4(1+w)} \Delta\tilde{\alpha}^R(x, \tilde{y}^R)$.

Proposition 6 shows that the same structure of the equilibrium solution from the baseline model continues to hold in this extended model. In particular, for any $0 < \sigma_L \leq \sigma_H \leq 1$, we can prove the similar properties on the shape of the $\tilde{M}(x)$ function, which ensure a single crossing of $\tilde{M}(x)$ and $2cx$, and hence establish the existence and uniqueness of $\tilde{x}^R$.

Proposition 6 lays the foundations for the robustness of the main results in subsection 4.4 when extended here. Next, we compare the equilibrium effort $\tilde{x}^R$ and $\tilde{y}^R$ and the service effectiveness, measured by $EV(\tilde{x}^R, \tilde{y}^R)$, with their counterparts when V is observed as a signal. We summarize all the results in Table 1, and discuss the key findings in more detail afterward.

Table 1 Equilibrium Comparison When Review Contains Private Information ($c > \frac{1}{8}$)

	Uninformative private signal $\sigma_H = \sigma_L = \sigma$ ($\sigma \in (0, 1]$)	Informative symmetric private signal $\sigma_H = \gamma$, $\sigma_L = 1 - \gamma$ ($\gamma \in [\frac{1}{2}, 1]$)
y^*	$y^V \geq \tilde{y}^R \geq y^e$ for any $0 < \sigma_L \leq \sigma_H \leq 1$ (The last inequality holds strictly if and only if $\sigma_H > \sigma_L$)	
x^*	(i) $\tilde{x}^R$ is increasing in σ; (ii) $\tilde{x}^R < x^V$, $\forall \sigma \in (0, 1]$	(i) $\tilde{x}^R$ is increasing in γ; (ii) $\tilde{x}^R > x^V$ if $\mu > \tilde{\mu}$ and $\gamma > \tilde{\gamma}$;† $\tilde{x}^R \leq x^V$ otherwise
$EV(x^*, y^*)$	$EV(\tilde{x}^R, \tilde{y}^R) < EV(x^V, y^V)$, $\forall \sigma \in (0, 1]$	$EV(\tilde{x}^R, \tilde{y}^R) > EV(x^V, y^V)$ if $w < \hat{w}$, $\mu > \hat{\mu}$, and $\gamma > \hat{\gamma}$;‡ $EV(\tilde{x}^R, \tilde{y}^R) \leq EV(x^V, y^V)$ otherwise

Notes: †There exist such thresholds $\tilde{\mu}(w, c) \in (0, 1)$ and $\tilde{\gamma}(\mu, w, c) \in (\frac{1}{2}, 1)$.
‡There exist such thresholds $\hat{w}(c) \in (0, 1)$, $\hat{\mu}(w, c) \in (0, 1)$ and $\hat{\gamma}(\mu, w, c) \in (\frac{1}{2}, 1)$.

PROPOSITION 7. *Comparing the equilibrium effort and the service effectiveness $\tilde{x}^R$, $\tilde{y}^R$, and $EV(\tilde{x}^R, \tilde{y}^R)$ with x^V, y^V, and $EV(x^V, y^V)$, we have the following results as summarized in Table* 1.

As Table 1 (the first row) shows, we first compare the equilibrium effort of the client ($\tilde{y}^R$ and y^V) under general values of σ_H and σ_L. The result that the review as a performance signal leads to a lower, yet more efficient, effort level for the client remains robust. In particular, when $\sigma_H = \sigma_L$, $\tilde{y}^R$ reaches the efficient level of effort y^e, which is consistent with the result under the baseline model (i.e., Proposition 5). Interestingly, when the client's private signal becomes informative (i.e., $\sigma_H > \sigma_L$), we have $\tilde{y}^R > y^e$. With an informative private signal, the client's review is more likely to reflect the true type of the provider and less sensitive to the client's utility. It hence reduces the incentive of the provider to cater to the client by prescribing a lower effort level in exchange for higher satisfaction. As a result, the provider requires the client to put in more effort than the most efficient level for him. Nevertheless, as long as the provider is still concerned about the client's utility when the review serves as a performance signal, $\tilde{y}^R$ remains more efficient than y^V in general.

We next compare the equilibrium efforts of the provider ($\tilde{x}^R$ and x^V) in two special cases, which allow us to derive formal results with representative and generalizable implications. The first case is when the client's private signal is not informative so that $\sigma_H = \sigma_L \equiv \sigma$, which subsumes the baseline model as $\sigma = 1$. The second case is when the client's private signal is informative and symmetric so that $\sigma_H = \gamma$ and $\sigma_L = 1 - \gamma$, $\gamma \in [\frac{1}{2}, 1]$. Recall that $\sigma_\theta = \Pr\{\tilde{\theta} = H|\theta\}$. In this sense, the symmetric structure indicates that with probability γ, the client's private signal correctly reflects the true type of the service provider. Note that $\gamma \geq \frac{1}{2}$ so that $\sigma_H \geq \sigma_L$, meaning that the client's private signal is informative or "correct" to a large extent.

As Table 1 (the second row) shows, the result that the review as a performance signal leads to less equilibrium effort of the provider continues to hold in general: $\tilde{x}^R < x^V$ holds for all σ's in the first case, and in most regions in the second case. An interesting aspect worth highlighting is that $\tilde{x}^R$ increases with the precision of the client's private signal γ, and once γ is large enough, it is possible that $\tilde{x}^R$ exceeds x^V. This new result can be well understood based on the interaction of the two effects that govern the provider's incentive of effort investment. As the client's private signal becomes more precise (i.e., as γ increases), the affirmative effect is partially weakened because a low type will be less likely to get a positive review anyhow, even if the outcome is successful. On the contrary, with precise private observation embedded in the review, a positive review becomes highly informative about the provider type. Hence, the informative effect significantly increases. As γ becomes large enough, the increase of the informative effect dominates the reduction of the affirmative effect, leading to higher marginal benefit for the provider's effort investment and hence higher equilibrium effort.

Recall that in the baseline model, the review brings down both the affirmative and the informative effects along the same direction. In contrast, when the review contains additional information about the provider type, it may change the two effects in the opposite directions, which drives the relative magnitude of the equilibrium effort one way or another. It is also worth emphasizing that because $\tilde{y}^R \leq y^V$ holds for any $\sigma_L \leq \sigma_H$, when the review is highly informative, it could be the case that $\tilde{x}^R$ compares to x^V in the opposite direction as $\tilde{y}^R$ compares to y^V. In other words, with the review rather than the outcome as the performance signal, the provider strategically lowers the effort level required for the client, but she works more diligently herself in compensation. In this sense, the collaborative effort decisions could exhibit a strategic-substitute pattern when the client review is highly informative, which enriches the strategic-complement pattern found under the baseline model.

We further compare the service effectiveness, $EV(\tilde{x}^R, \tilde{y}^R)$ and $EV(x^V, y^V)$. As Table 1 (the third row) shows, we find the similar result that the review generally leads to less effectiveness of service compared to the outcome, that is, $EV(\tilde{x}^R, \tilde{y}^R) < EV(x^V, y^V)$ in general. Nevertheless, such a relationship can be reversed when the review contains highly precise private observation of the client (i.e., in the second case with γ sufficiently large), just like the comparison between $\tilde{x}^R$ and x^V as discussed above. As Figure 2a (the dark grey region) plots, when w is small (i.e., the client's effort is not too costly) and μ is large (i.e., the ability gap between the two types of providers is small enough), $EV(\tilde{x}^R, \tilde{y}^R)$ could exceed $EV(x^V, y^V)$ for large γ's.

To demonstrate that the results derived from the two special cases shown in Table 1 are representative and general, we further solve the equilibrium numerically for all different values of σ_H and σ_L ($0 < \sigma_L \leq \sigma_H \leq 1$). Figure 2b depicts the comparison result of $\tilde{x}^R$ versus x^V over the plane of σ_H and σ_L. Notice that the two special cases correspond to the two diagonals. When moving from the center towards the lower right corner of the figure, σ_H increases and σ_L decreases, representing that the client's private signal becomes more and more informative or accurate. As can be clearly seen from Figure 2b, $\tilde{x}^R$ exceeds x^V only around the lower right corner, which is consistent with the implications from the two special cases in Table 1.

5.2. Outcome and Review as Signals

Under this extended model, when both the outcome V and the review R are observable to the market as performance signals, the equilibrium solution will be different from when only V is available. The reason is that the conditional independence of R on θ given V as in (4) no longer holds under this extended model. According to (15), the probability of receiving a positive or negative review, even given the outcome V, still depends on the provider type θ. Thus, $\Pr\{V, R|\theta, x, y\} = \Pr\{R|V, \theta, y\}\cdot\Pr\{V|\theta, x, y\}$. Substitute this likelihood function into (6), and we can easily see that the posterior belief does depend on the realized values of both V and R.

We can follow the same approach described previously to derive the equilibrium solution here. The only challenge lies in that the first-order conditions no longer separate the optimal solution of y from the optimal solution of x as in the previous scenarios. Instead, we need to solve two simultaneous equations of higher order involving both x and y, which yields no closed-form analytical solution. Nevertheless, we can solve the equilibrium numerically under different parameter values. Figure 3 illustrates the equilibrium efforts, $\tilde{x}^{VR}$ and $\tilde{y}^{VR}$, in the case of symmetric informative signal such that $\sigma_H = \gamma$ and $\sigma_L = 1 - \gamma$ ($\gamma \in [\frac{1}{2}, 1]$), in comparison to the equilibrium efforts under the other two scenarios. Consistently, we use the superscript VR to denote the equilibrium solution when both V and R are available as signals.

As Figures 3a and b show, when γ is relatively small, $\tilde{y}^{VR}$ takes the corner solution so that $\tilde{y}^{VR} = 1 = y^V$ and $\tilde{x}^{VR} = x^V$. The most surprising result revealed in Figures 3a and b is that both $\tilde{x}^{VR}$ and $\tilde{y}^{VR}$*decrease* as γ becomes larger, changing in the opposite direction as $\tilde{x}^R$ and $\tilde{y}^R$. In other words, when both the outcome and the review are observable by the market, as the review itself is more informative about the true provider type, there is *less* incentive for the service provider to devote effort herself and to require the client to work diligently—in fact,

Figure 2 Equilibrium Comparison When Review Contains Private Information (c = 0.15)

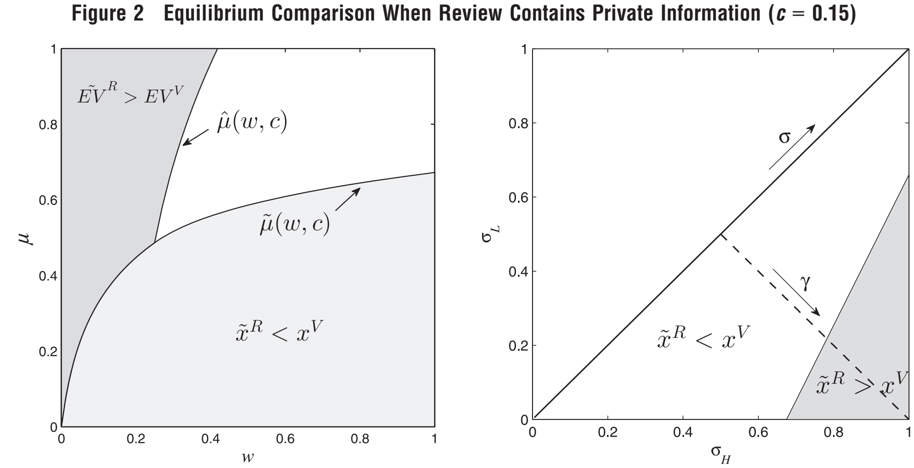

(a) Equilibrium Comparison with $\gamma = 1$

(b) Equilibrium Comparison with General σ_H and σ_L (w = 0.25, μ = 0.7)

Figure 3 Equilibrium Comparison with $\sigma_H = \gamma$ and $\sigma_L = 1 - \gamma$ ($c = 0.15$, $w = 0.25$, $\mu = 0.7$)

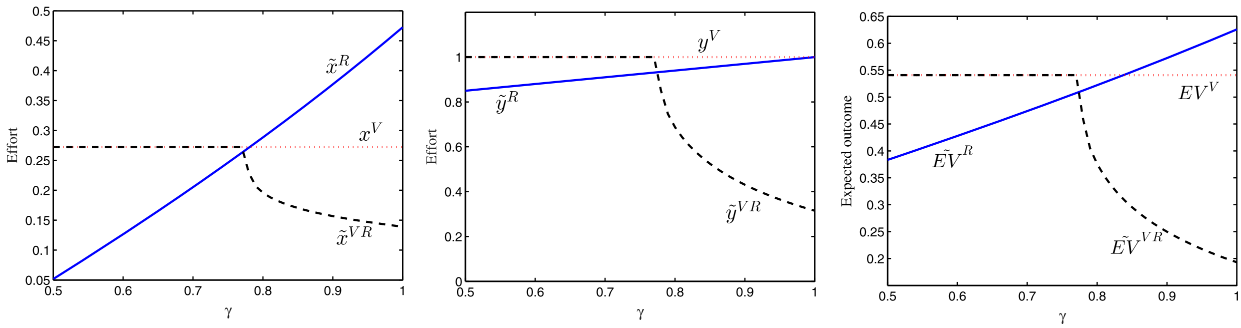

(a) Equilibrium Effort of the Provider **(b)** Equilibrium Effort of the Client **(c)** Effectiveness of Service

achieving a successful outcome is less desirable for the service provider in this case. This counterintuitive result originates from the non-monotonicity of the informative effect under a bivariate signal structure. When the market signal is univariate (i.e., either V or R alone is observed), a positive signal is always positively informative in that it suggests the service provider is more likely to be a high type than a low type; however, with a bivariate signal structure (i.e., both V and R are observed), this may not always be the case. For example, if the service provider turns out to be a low type, with a highly accurate private signal of the client (i.e, γ close to 1), the provider will most likely receive a negative review; in this case, the provider might be even better off with a failed outcome than a successful outcome. This is because the combination of $\{V = 1, R = 0\}$ could convey a stronger message that the provider is a low type than the combination of $\{V = 0, R = 0\}$, considering that receiving a negative review despite a successful outcome suggests to the market that the client might have observed a negative private signal, which is supposedly accurate. Because of such non-monotonicity, the overall expected benefit of achieving a successful outcome through costly effort investment significantly reduces when γ increases under the bivariate signal structure, resulting in the interesting decreasing pattern of both $\tilde{x}^{VR}$ and $\tilde{y}^{VR}$ found in Figures 3a and b.

Given that the equilibrium effort levels $\tilde{x}^{VR}$ and $\tilde{y}^{VR}$ both decrease as the review becomes more informative, it can be expected that the effectiveness of service, measured as $EV(\tilde{x}^{VR}, \tilde{y}^{VR})$, also decreases in γ when γ is large, as we show in Figure 3c. In fact, when the review contains accurate information about the provider type, the scenario with both the outcome and the review as signals leads to the lowest outcome expectation among all three scenarios.

We further plot the equilibrium judgment error $\tilde{\rho}^{VR}$, together with $\tilde{\rho}^{R}$ and ρ^{V} in Figure 4, where 4a shows how they change in μ, and 4b in γ. Surprisingly, contrary to the intuition that more information typically leads to better distinguishability, as Figure 4 shows, with two pieces of information observable (i.e., both V and R), $\tilde{\rho}^{VR}$ could end up higher than $\tilde{\rho}^{R}$, the equilibrium judgment error when only one piece of information is observable. The reason is the lower effort levels strategically chosen by the service provider in equilibrium, as discussed previously. This interesting result thus serves as a good illustration of how strategic interplay could garble the information effectively conveyed under a complex information structure.

As we can also see in Figure 4, ρ^{V} is low when μ is small, indicating that the outcome leads to a high degree of distinguishability only when the inherent ability gap between the two types of providers is large. Recall the results in Figure 1 that the review generally leads to a greater error. In contrast, as shown in Figure 4, with additional information about the provider type, the review alone can result in the lowest equilibrium judgment error among all three scenarios, especially when μ and γ are large. Along this line, summarizing the analysis in this section sheds light on the effects of client reviews on service performance: with reasonably accurate extra information about the provider type, client reviews could lead to superior performance of collaborative services with better distinguishability of provider types, greater service effectiveness, and relatively more efficient effort requirement for the client.

6. Extensions

6.1. Client's Effort Jointly Determined

In this section, we extend the model to relax the assumption that the service provider solely determines the effort level of the client in the collaborative service. We now consider the general case that both the provider and the client jointly determine the effort level of client, and we allow the two parties to have their respective degrees of influence, which can vary as a model parameter. Such a generalization thus subsumes the two special cases when either the provider or the client alone fully determines the client's effort.

Figure 4 **Comparison of Equilibrium Judgment Errors (c = 0.15, w = 0.25)**

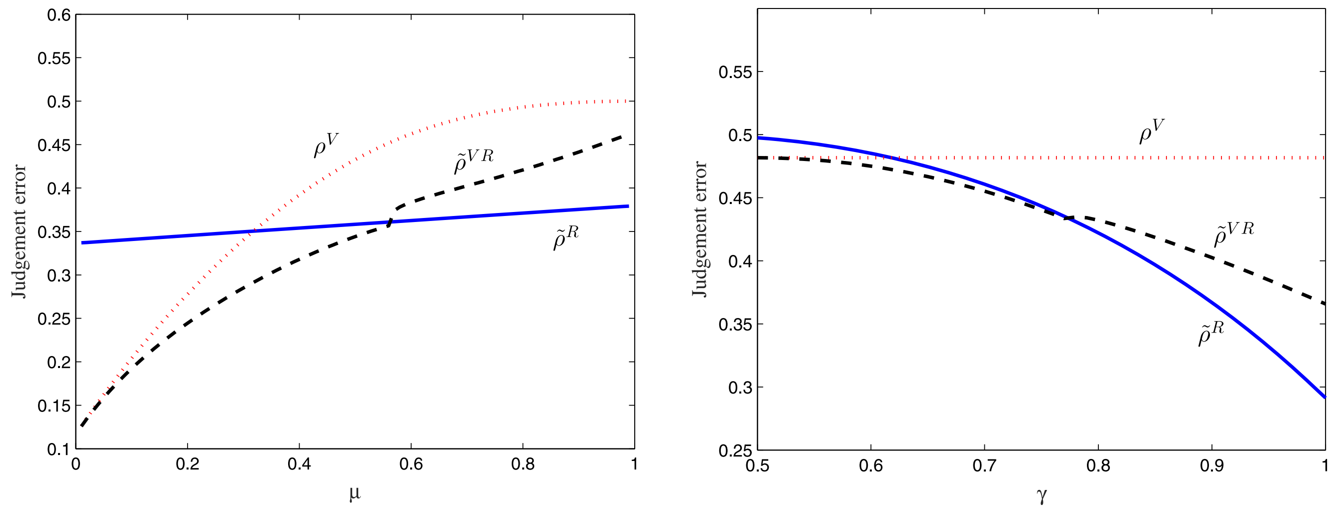

(a) ρ^* Changes in μ ($\gamma = 0.9$)

(b) ρ^* Changes in γ ($\mu = 0.7$)

Specifically, we let the overall effort level the client eventually exerts, y, be jointly determined by a level required by the service provider, y_1, and a level chosen by the client himself, y_2, as follows.

$$y = \lambda y_1 + (1-\lambda)y_2, \quad \lambda, y_1, y_2 \in [0,1]. \tag{16}$$

Here, $\lambda \in [0, 1]$ captures the relative strength of influence of the two parties. If $\lambda = 1$, $y = y_1$, so the service provider fully determines y, and it hence reduces to the original model. When $\lambda = 0$, $y = y_2$, so it is totally up to the client to decide his own effort, which is another special case of our interest. The service provider first decides y_1, and the client chooses y_2 after observing the required level y_1. As previously, both y_1 and y_2 are decided and invested upfront before the outcome realizes. To differentiate from the previous sections, we add a hat ($\wedge$) to indicate the equilibrium solutions in this extended model.

In determining his optimal effort choice $\hat{y}_2^*$, the client maximizes the *ex ante* expectation of his net utility $U(V, y) = \mathbb{1}_{V=1} - wy^2$, and therefore,

$$\hat{y}_2^*(y_1) = \arg\max_{0 \le y_2 \le 1} E_\theta\{E[V|\theta, \hat{x}^*, \lambda y_1 + (1-\lambda)y_2]\} - w(\lambda y_1 + (1-\lambda)y_2)^2. \tag{17}$$

The service provider sets the optimal $\hat{y}_1^*$ by maximizing her expected payoff according to (7), where $y = \lambda y_1 + (1-\lambda)\hat{y}_2^*(y_1)$. Notice that the optimal overall effort level y for the client (if he could fully decide) would be the efficient effort level y^e as defined by (12); the optimal y for the service provider (if she could fully decide) would simply be what is derived from the original model, y^*. Hence, the resulting overall effort level of the client in this extended model, $\hat{y}^*$, is a balance between y^* and y^e, depending on the two parties' relative strength of influence, λ. When the provider has strong influence, for example, if $\lambda > y^* > y^e$, the provider will set $\hat{y}_1^* = y^*/\lambda$, anticipating that the client will try to lower the overall effort as close to y^e as possible by setting $\hat{y}_2^* = 0$; consequently, $\hat{y}^* = y^*$. When the client has strong influence, for example, if $\lambda < y^e < y^*$, the provider could actually set $\hat{y}_1^*$ as any value between $\max\{\frac{y^e+\lambda-1}{\lambda}, 0\}$ and 1, because the client will choose $\hat{y}_2^* = \frac{1}{1-\lambda}(y^e - \lambda\hat{y}_1^*)$ to always bring the overall effort level $\hat{y}^*$ equal to y^e eventually. In other cases, for example, if $y^e < \lambda < y^*$, $\hat{y}^*$ will simply be equal to λ as a result of $\hat{y}_1^* = 1$ and $\hat{y}_2^* = 0$.

We first extend the baseline model to incorporate the jointly determined effort structure. The equilibrium effort when either the outcome or the review is observable as a signal is summarized in Table 2 (the first two columns). When the review is observed as a signal, recall that Proposition 5 shows the optimal y chosen by the service provider, y^R , turns out to be exactly the same as the efficient effort level the client himself would choose, y^e. In this case, the provider's and the client's strategies completely align. Therefore, in this extended model, regardless of the value of λ, the overall level of the client's effort remains the same as the one derived from the baseline model, that is, $\hat{y}^R = y^R = y^e = \min\{\frac{1+\mu}{8w}, 1\}$. When the outcome is observed as a signal, however, the provider's and the client's optimal choices of y differ, and $\hat{y}^V$ depends on the value of λ. If $\lambda = 1$, $\hat{y}^V$ reduces to a constant 1, which equals y^V in the baseline model. If $\lambda = 0$, the

Table 2 Equilibrium Effort under Jointly Determined *y*

V as a signal	R as a signal (w/o private information)	R as a signal (w/ private information, $\sigma_H = \gamma$, $\sigma_L = 1 - \gamma$)
$\hat{y}^V = \begin{cases} 1 & \text{if } \frac{1+\mu}{8w} \geq 1 \\ \frac{1+\mu}{8w} & \text{if } \lambda < \frac{1+\mu}{8w} < 1 \\ \lambda & \text{if } \frac{1+\mu}{8w} \leq \lambda \end{cases}$	$\hat{y}^R = \begin{cases} 1 & \text{if } \frac{1+\mu}{8w} \geq 1 \\ \frac{1+\mu}{8w} & \text{if } \frac{1+\mu}{8w} < 1 \end{cases}$	$\hat{\tilde{y}}^R = \begin{cases} 1 & \text{if } w \leq \frac{1+\mu}{8} \\ \frac{1+\mu}{8w} & \text{if } \frac{1+\mu}{8} < w \leq \frac{1+\mu}{8\lambda} \\ \lambda & \text{if } \frac{1+\mu}{8\lambda} < w \leq \frac{\gamma+(1-\gamma)\mu}{4\lambda} \\ \frac{\gamma+(1-\gamma)\mu}{4w} & \text{if } w > \frac{\gamma+(1-\gamma)\mu}{4\lambda} \end{cases}$
$\hat{x}^*$ is the unique solution within [0, 1] to the following equations ($c > 1/8$)		
$K(\hat{x}^V, \hat{y}^V) = 2c\hat{x}^V$	$M(\hat{x}^R, \hat{y}^R) = 2c\hat{x}^R$	$\tilde{M}(\hat{\tilde{x}}^R, \hat{\tilde{y}}^R) = 2c\hat{\tilde{x}}^R$

where

$$K(x, y) = \frac{1-\mu}{2[4-(1+\mu)(x+y)]}.$$

$$M(x, y) = \left[\frac{(x+y)+2w(1-y^2)}{(1+\mu)(x+y)+4w(1-y^2)} - \frac{2(1+w)-(x+y)-2w(1-y^2)}{4(1+w)-(1+\mu)(x+y)-4w(1-y^2)}\right]\frac{1+\mu}{4(1+w)}.$$

$$\tilde{M}(x, y) = \left[\frac{\gamma(x+y)+2\gamma w(1-y^2)}{(\gamma+(1-\gamma)\mu)(x+y)+2w(1-y^2)} - \frac{2(1+w)-\gamma(x+y)-2\gamma w(1-y^2)}{4(1+w)-(\gamma+(1-\gamma)\mu)(x+y)-2w(1-y^2)}\right]\frac{\gamma+(1-\gamma)\mu}{4(1+w)}.$$

client himself fully determines his effort and hence chooses the efficient effort level. As a result, $\hat{y}^V$ reduces to $y^e (= \hat{y}^R = y^R)$. On the other hand, the equilibrium effort levels of the provider in this extended model (i.e., $\hat{x}^V$ and $\hat{x}^R$ characterized in Table 2) extend naturally from the baseline model. While the value of $\hat{x}^*$ could change as $\hat{y}^*$ differs from the baseline model in some cases, the functional relationship between $\hat{x}^*$ and $\hat{y}^*$ remains the same. Comparing $\hat{x}^V$ and $\hat{x}^R$ ($\hat{y}^V$ and $\hat{y}^R$) in this extended model, we can show that the same result derived from the baseline model continue to hold, as summarized by the following proposition.

PROPOSITION 8. *In the extended model with the client's effort jointly determined by both parties, for* $c > \frac{1}{8}$, $\hat{x}^R < \hat{x}^V$ *for any* $\lambda \in [0, 1]$, *and* $\hat{y}^R \leq \hat{y}^V$ *with strict inequality when* $\lambda > \frac{1+\mu}{8w}$.

We further apply the jointly determined effort structure to the case when the client incorporates symmetric private signal into the reivew, as discussed in Section 5. As shown in the rightmost column of Table 2, when $\lambda = 1$, $\hat{\tilde{y}}^R$ reduces to $\tilde{y}^R = \min\{\frac{\gamma+(1-\gamma)\mu}{4w}, 1\}$ as expected; when $\lambda = 0$, $\hat{\tilde{y}}^R$ reduces to $y^e = \min\{\frac{1+\mu}{8w}, 1\}$. Moreover, $\hat{\tilde{x}}^R$ extends naturally from Section 5. The results regarding the comparison between $\hat{\tilde{x}}^R$ and $\hat{x}^V$ ($\hat{\tilde{y}}^R$ and $\hat{y}^V$) continue to hold under this extended model, as follows.

PROPOSITION 9. *In the extended model with the client's effort jointly determined by both parties, when the client incorporates symmetric private information into the review (i.e.,* $\sigma_H = \gamma$ *and* $\sigma_L = 1 - \gamma$*), (i)* $\hat{\tilde{y}}^R \leq \hat{y}^V$ *with strict inequality when* $\lambda > \frac{\gamma+(1-\gamma)\mu}{4w}$*; (ii) for* $c > \frac{1}{8}$, *there exists* (c, w, μ, λ) *such that for a certain threshold* $\tilde{\gamma}(c, w, \mu, \lambda) \in (\frac{1}{2}, 1)$, $\hat{\tilde{x}}^R \leq \hat{x}^V$ *when* $\gamma \leq \tilde{\gamma}$, *and* $\hat{\tilde{x}}^R > \hat{x}^V$ *when* $\gamma > \tilde{\gamma}$.

In sum, the analysis in this section shows that when extended to allow the client to determine his effort input jointly with the service provider, our main results continue to hold. In the extreme case when the client solely determines his own effort (i.e., $\lambda = 0$), all the main results with respect to the provider's effort, $\hat{x}^*$, remain robust; meanwhile, the equilibrium effort of the client, $\hat{y}^*$, reduces to a simple case in which it always equals the efficient effort level, no matter what signal is observed by the market. In this sense, studying the general case that the provider (at least to some extent) decides the client's effort level in collaborative services allows us to further examine how different signals could affect the client's effort and to explore the additional strategic correlation between the two effort decisions.

6.2. Client's Effort Observable

Thus far, we consider the client's effort as unobservable to the market. There is another possible scenario that the client may convey information about his effort in his review, and as a result, the client's effort could be observable to the market when the review is available as a signal. In order to show that our model can accommodate such possibility and the main results remain robust, we conduct an extension along this direction. Due to page limitation, we include the detailed analysis and results in Appendix S4. We briefly discuss the intuition below. To differentiate from the original results, we add a dot on top of the equilibrium solutions under this extended model.

We follow the baseline model with the only difference that the client's effort y is now observed by the market when the review R serves as a signal. This difference leads to a major change in the equilibrium concept: Now that y is directly observed, the market incorporates this information in forming its rational belief about the provider type and "figuring out" the provider's equilibrium effort $\dot{x}^*$. Specifically, given any choice of y, a perfect Bayesian equilibrium is a pair $\{\dot{x}^*(y), \dot{\alpha}^*(y)\}$ such that $\dot{x}^*(y)$ is optimal for the provider given the market belief $\dot{\alpha}^*(y)$, while $\dot{\alpha}^*(y)$ is rational according to Bayes' rule given the provider's chosen effort $\dot{x}^*(y)$. Anticipating such equilibrium outcomes, the service provider chooses $\dot{y}^*$ that maximizes his expected payoff *ex ante*, which leads to his own equilibrium effort $\dot{x}^*(\dot{y}^*)$.

As we find, the equation defining $\dot{x}^R$ given y remains the same as previously, and therefore, $\dot{x}^R$ and $\dot{y}^R$ follow the same functional relationship. The provider's optimization of y results in $\dot{y}^R$ that minimizes her own effort input $\dot{x}^R(y)$ among all possible y's. As a result, $\dot{x}^R(\dot{y}^R) \leq \dot{x}^R(y^R) = x^R$, that is, the equilibrium effort of the provider in this extended model $\dot{x}^R$ is no greater than that in the original model x^R. Notice that the equilibrium when the outcome V serves as a signal is unaffected in this extension. Comparing $\dot{x}^R$ and x^V ($\dot{y}^R$ and y^V), we can show that $\dot{y}^R \leq y^V(\equiv 1)$ and $\dot{x}^R(\leq x^R) < x^V$ continue to hold.

We further extend the case when the client incorporates private information into the review to allow the possibility of y being observable. As we find, similar results continue to hold: when the client's private information is precise enough (i.e., when γ is sufficiently large), it is possible that $\dot{\hat{x}}^R > x^V$, whereas $\dot{\hat{y}}^R \leq y^V$ always holds, resulting in a strategic-substitute pattern between the two effort decisions.

7. Conclusion

In this study, we take the first step to study the provision of collaborative services under online reviews. Utilizing a signal-jamming model, we endogenously analyze how effectively the client review conveys information about the underlying capability of the service provider and examine the interplay between such endogenous informativeness and the strategic actions of the provider. Our modeling results provide rich managerial implications for information systems designers, top management and authorities, service providers, and clients as well.

A key implication to information systems designers is that the information being collected and made available changes the strategic behavior of system users and hence impacts the performance the systems aim to achieve, especially in the context of service management. In fact, the common wisdom that more information is better may not always hold. As we show, when both the client review and the service outcome are observable to the public, the motivation for service providers to devote effort could be unexpectedly hurt. As a result, two pieces of information may lead to lower effort levels and hence inferior performance (i.e., effectiveness of service, distinguishability of provider types) than one piece of information alone. Therefore, our findings underscore the crucial importance of endogenously examining the strategic consequences of the proposed information structure in designing information systems for service management.

For top management and authorities (e.g., managers of private tutoring companies, administrators of a school district) to deploy evaluation systems to incentivize subordinate service providers to improve their service provision, our study emphasizes that such decisions need to be context specific and could differ greatly depending on the service characteristics. In general, two dimensions of the service characteristics are particularly important to consider: the ability difference among service providers, and how accurately clients can tell the difference among providers after the service. Our results show that reviews serve as a satisfactory performance indicator when the difference among service providers are not too large (e.g., for industries when all service providers need to pass relatively strict qualification or certification processes); otherwise, service outcomes would be a better choice in general. For services in which clients are able to relatively accurately tell the difference among providers (e.g., mature adults, experienced learners), client reviews could lead to superior effects in multiple dimensions; in contrast, the effects of reviews may be significantly impaired if, for example, the clients are mainly young children.

Our modeling results provide further guidelines for collaborative service providers in determining their effort strategies. We suggest that effort input is the most rewarding when reviews are available and clients incorporate accurate private observations into their reviews. In this case, it is wise for service providers to refrain from pushing clients to overwork and meanwhile to compensate with excessive effort themselves. Finally, our study also provides implications for clients who review their service providers. As we show, reviews with more informative private signals generally lead to more desirable results in various dimensions. In this sense, as clients write reviews, whenever possible, they should also rely on their personal observations about the provider's capability than merely their utility based on the service outcome and their effort.

As an initial attempt, this study is not without limitations, which point to interesting directions for future research. First, we focus only on the services in which the service outcome is beyond the manipulation of the service provider. Once the provider can strategically manipulate the service outcome perceived by the client, this additional strategic decision will naturally lead to interesting new findings. Second, we do not model the pricing decision of the service provider. As the motivating example of our study, the industry of educational and tutoring services consists of mostly small service providers, and there is usually a prevailing price rate for providing a certain type of service in each particular geographic area (Franchisehelp 2016). Therefore, service providers typically act as price takers. Besides, in many occasions, the clients who write reviews do not incur the service cost directly,[7] so the price of the service does not enter the client's utility function to affect the review. Based on these considerations, we choose to leave the pricing of the service as exogenous in order to focus on the strategic interplay between the effort levels chosen for both the provider and the client. Nevertheless, allowing the provider to price endogenously will be an important and interesting future research direction. Third, we show that a bivariate signal structure may lead to counterintuitive equilibrium outcomes through numerical analysis. A general modeling study on the strategic interplay under multivariate signal structure in a signal-jamming game would thus be a challenging yet intriguing direction with theoretical significance to information economics.

Acknowledgments

The authors contribute equally and are listed in alphabetical order. Lizhen Xu is the corresponding author. The authors thank Dr. Vijay Mookerjee, the senior editor, and two anonymous reviewers for their insightful and constructive comments, which have greatly improved this study. The authors also thank the session participants at INFORMS Annual Meeting (2013) and POMS Annual Conference (2016) for their helpful feedback.

Notes

[1]In this study, we focus on the case in which the outcome of the service is objective and beyond the manipulation of service providers (e.g., passing a standardized test, achieving preset goals). Other cases in which service providers may be able to manipulate the service outcome (e.g., by lowering the difficulty of the test) involve additional strategic interplay beyond the scope of this study and are therefore left for future research.

[2]The client can be viewed as the representative of a group. Our model can also be extended to the case of a finite number of clients, and the main results will continue to hold qualitatively.

[3]For example, Wherry and Bartlett (1982) develop a psychological theory decomposing rating into a systematic component and a random component. Kane (1994) identifies multiple idiosyncratic factors (e.g., halo effect, assimilation or contrast effects, primacy or recency effects) that could bias ratings. Li and Hitt (2010) find empirical evidence that customer ratings are correlated more closely with the net utility that customers derive from the product than with product quality itself. Net utility-based rating is commonly adopted in theoretical work (e.g., Jiang and Guo 2015, Kuksov and Xie 2010).

[4]Our results can be further extended to the case when y is observable to the market along with online reviews, as we discuss in subsection 6.2.

[5]It would also be interesting to examine the efficiency of the provider's own effort choice. Nevertheless, the definition of the efficient effort level of the service provider is indecisive in our context. It depends on how to interpret the market rewards so as to define the overall social welfare. The result will also depend on the values of m_H and m_L (which are normalized for simplicity) relative to other model parameters. For these reasons, we choose to leave out this part of analysis to avoid sidetracked discussions.

[6]If the prior belief is not symmetric, the equilibrium type I and type II errors themselves may differ. However, the general trend regarding the comparison result across the two scenarios will remain the same.

[7]For example, in professional training or consulting projects, while companies pay for the services, their employees deal with the details and evaluate the service providers. Similar examples exist in the education industry: for instance, U.S. federal government established the Supplemental Educational Services (SES) to provide free after-school tutoring to low-income students who struggle academically. Through SES, parents choose any private providers on the state's vendor lists and the school districts pay the costs of tutoring.

References

Bakos, Y., C. Dellarocas. 2011. Cooperation without enforcement? A comparative analysis of litigation and online reputation as quality assurance mechanisms. *Management Sci.* **57**(11): 1944–1962.

Bhattacharya, S., A. Gupta, S. Hasija. 2014. Joint product improvement by client and customer support center: The role of gain-share contracts in coordination. *Inf. Syst. Res.* **25**(1): 137–151.

Chen, Y., J. Xie. 2005. Third-party product review and firm marketing strategy. *Mark. Sci.* **24**(2): 218–240.

Chen, Y., J. Xie. 2008. Online consumer review: Word-of-mouth as a new element of marketing communication mix. *Management Sci.* **54**(3): 477–491.

Chetty, R., J. N. Friedman, J. E. Rockoff. 2014. Measuring the impacts of teachers II: Teacher value-added and student outcomes in adulthood. *Am. Econ. Rev.* **104**(9): 2633–2679.

Chevalier, J. A., D. Mayzlin. 2006. The effect of word of mouth on sales: Online book reviews. *J. Mark. Res.* **43**(3): 345–354.

Demirezen, E., S. Kumar, B. Shetty. 2013. Value co-creation in IT outsourcing and software development. *Conference on Information Systems and Technology*.

Franchisehelp. 2016. Education industry analysis 2016—cost & trends. Market research report, FranchiseHelp Holdings LLC. Available at https://www.franchisehelp.com/industry-reports/educational-franchise-industry-report/ (accessed date: January 16, 2016).

Fudenberg, D., J. Tirole. 1986. A "signal-jamming" theory of predation. *RAND J. Econ.* **17**(3): 366–376.

GIA. 2014. Private tutoring—a global strategic business report. Market research report, Global Industry Analysts, Inc. Available at http://www.strategyr.com/pressMCP-1597.asp (accessed date: March 25, 2015).

Gu, B., J. Park, P. Konana. 2012. The impact of external word-of-mouth sources on retailer sales of high-involvement products. *Inf. Syst. Res.* **23**(1): 182–196.

Hao, L., X. Li, Y. Tan, J. Xu. 2011. The economic value of ratings in app market. Working paper, Available at http://ssrn.com/abstract=1892584 (accessed date: March 25, 2015).

Holmstrom, B. 1999. Managerial incentive problems: A dynamic perspective. *Rev. Econ. Stud*. **66**(1): 169–182.

Jiang, Y., H. Guo. 2015. Design of consumer review systems and product pricing. *Inf. Syst. Res.* **26**(4): 714–730.

Kane, J. S. 1994. A model of volitional rating behavior. *Hum. Resou. Manag. Rev.* **4**(3): 283–310.

Karmarkar, U. S., R. Pitbladdo. 1995. Service markets and competition. *J. Oper. Manag.* **12**(3–4): 397–411.

Kreps, D. M., R. Wilson. 1982. Reputation and imperfect information. *J. Econ. Theory* **27**(2): 253–279.

Kuksov, D., Y. Xie. 2010. Pricing, frills, and customer ratings. *Mark. Sci.* **29**(5) 925–943.

Kwark, Y., J. Chen, S. Raghunathan. 2014. Online product reviews: Implications for retailers and competing manufacturers. *Inf. Syst. Res.* **25**(1): 93–110.

Li, X., L. M. Hitt. 2010. Price effects in online product reviews: An analytical model and empirical analysis. *MIS Q.* **34**(4): 809–831.

Liu, Y. 2006. Word of mouth for movies: Its dynamics and impact on box office revenue. *J. Mark.* **70**(3): 74–89.

Milgrom, P., J. Roberts. 1982. Predation, reputation, and entry deterrence. *J. Econ. Theory* **27**(2): 280–312.

Plambeck, E. L., T. A. Taylor. 2006. Partnership in a dynamic production system with unobservable actions and noncontractible output. *Management Sci.* **52**(10): 1509–1527.

Roels, G., U. S. Karmarkar, S. Carr. 2010. Contracting for collaborative services. *Management Sci.* **56**(5): 849–863.

Spohrer, J., P. P. Maglio. 2008. The emergence of service science: Toward systematic service innovations to accelerate co-creation of value. *Prod. Oper. Manag.* **17**(3) 238–246.

Sun, M. 2012. How does the variance of product ratings matter? *Management Sci.* **58**(4): 696–707.

Wherry, R. J.Sr., C. J. Bartlett. 1982. The control of bias in ratings: A theory of rating. *Pers. Psychol.* **35**(3): 521–551.

Xue, M., J. M. Field. 2008. Service coproduction with information stickiness and incomplete contracts: Implications for consulting services design. *Prod. Oper. Manag.* **17**(3): 357–372.

Zhu, F., X. M. Zhang. 2010. Impact of online consumer reviews on sales: The moderating role of product and consumer characteristics. *J. Mark.* **74**(2): 133–148.

Supporting Information

Additional Supporting Information may be found in the online version of this article:

Appendix S1: Proofs.
Appendix S2: Different Distributions for ξ and ε.
Appendix S3: Results When c is Small.
Appendix S4: Results for Sections 6.2.