Implications of Social Preferences in a Services Market with Heterogeneous, Uninformed Consumers

Jian Ni[*]
jni@jhu.edu, Carey Business School, Johns Hopkins University

Baojun Jiang
baojunjiang@wustl.edu, Olin Business School, Washington University in St. Louis

Kannan Srinivasan
kannans@cmu.edu, Tepper School of Business, Carnegie Mellon University

September 2012

Abstract

In many services markets such as consulting and healthcare, the service provider may have more information about the client's problem than the client, and different clients may impose different costs on the service provider. In principle, the service provider should ethically care about the client's welfare, but it is possible that a provider may maximize only its own profit. Moreover, the client may not know ex ante whether the service provider is ethical or purely self-interested. Using healthcare as the motivating context, we develop a game-theoretic model to investigate pricing strategies and the market outcome in a services market where the service provider (i.e., the doctor) has two-dimensional private information—about her own type (whether ethical or self-interested) and about the patient's condition (whether severe or minor). We show that, in a less ethical market, where the probability of a doctor being ethical is low, a self-interested doctor will charge different prices based on the severity of the patient's condition whereas an ethical doctor will charge the same price for both severity conditions. In contrast, in a more ethical market, both the self-interested and the ethical doctor will charge the same uniform price to both types of patients. Interestingly, the market efficiency may be lower in a more ethical market than in a less ethical one, and patients with minor conditions become worse off and those with severe conditions become better off as the probability of the doctor being ethical increases.

Keywords: social preference, signaling, pricing, behavioral economics, asymmetric information

---

[*] All authors contributed equally and are listed in a random order.

1. Introduction

Some services markets such as consulting, investment management, and healthcare exhibit three characteristics. First, in many such markets, the service provider may have more information about the client's problem than the client. The consulting firm may have extensive expertise in dealing with specific technical, management, or strategic issues because it may have dealt with similar issues with many other clients and acquired much more information to be able to provide accurate diagnosis and satisfactory resolution. Similarly, the doctor or healthcare services provider typically knows much more about a patient's medical condition than the patient himself.[1]

Second, different clients' problems may impose different levels of cost on the service provider and the clients may not know the provider's true costs for resolving an issue. This is very different from the traditional product market, where the product's cost is usually the same across consumers. In the context of services, the provider's cost may be customer-specific. Different conditions require different service/effort levels and consequently have different costs to the firms. This, as we show later, might lead the firm to dump some customers because of the high service cost.

Third, even though in principle the service provider should ethically care about the client's welfare (e.g., by fiduciary duty or ethical codes of conduct), it is possible that a provider may seek only to maximize its own profit. For example, a doctor might not always heed the welfare of her patients. Some healthcare services providers prescribe more than adequate or necessary levels of care to patients (Jaegher and Jegers, 2000). Some even stop providing procedures or services that carry a high liability of risk or require a great amount of time and effort, or simply dump those patients with severe conditions that cost them much more to service (Ansell and Schiff, 1987). A survey of 1900

---

[1] For expositional ease, we will refer to a patient as "he" and a doctor as "she."

physicians in the U.S. indicated that they do not always tell the truth, e.g., a third of the doctors surveyed did not think that they should disclose medical errors and 40% of them did not think they needed to disclose their financial ties to drug and devices companies.[2] Another example is that during the global financial crisis that started in 2007, some investment management companies acted against the clients' interests and did not disclose potential conflicts of interest to clients when recommending and carrying out financial transactions, even though their standards state that their clients' interests always come first. For example, the former Goldman Sachs executive, Greg Smith, said in his open resignation letter in the *New York Times* that "the interests of the client continue to be sidelined in the way the firm operates and thinks about making money."[3] Clearly, some service providers, be it doctors or financial services companies or managers, do not always heed the welfare of their clients as they should according to ethical standards or business principles. Moreover, the clients may not know ex ante whether a service provider is ethical or purely self-interested in any business or health services interactions.

Using healthcare as the motivating and expositional context, we develop a game-theoretic model to investigate the pricing strategies and market outcomes in a services market with the aforementioned key characteristics where the doctor (i.e., the service provider) has two-dimensional private information—one about her own type (whether ethical or self-interested) and the other about the patient's medical condition (whether severe or minor). Patients often know neither which health services they need nor to what extent the healthcare services that they have received contributed to their health. This informational asymmetry induces inefficient patient search

---

[2] See http://healthland.time.com/2012/02/09/white-coats-white-lies-how-honest-is-your-doctor/, accessed in April 2012.

[3] See http://www.nytimes.com/2012/03/14/opinion/why-i-am-leaving-goldman-sachs.html?pagewanted=all, accessed in April 2012.

(Wolinsky, 1993) and sometimes leads to patients' rejecting necessary treatments (Fong, 2005). Both cases can result in significant welfare loss.

Even though the quality of provider-patient interactions has been declining, the medical profession's ethical codes enforce a policy of self-restraint (Farley, 1986; Kaufman, Fein, and Fins, 2009). Standard economic literature generally assumes that firms maximize their monetary profits. But the existence of social preference such as conscientious or altruistic behavior is demonstrated by extensive behavior literature. For example, ethical doctors may derive some utility from not only their own profits but also their patients' welfare (e.g., Ho et al., 2006a, 2006b). This explains why some doctors keep their promises to treat their patients even though the associated costs may be higher than the price they charge (Vanberg, 2008).

Unfortunately, a self-interested doctor may mimic an ethical doctor when the patient's problem is minor or more profitable to the doctor, but reject the severe-condition or less profitable patients unless they can charge a high enough price. Further, the doctor may charge a patient a high price for minor (low-cost and less serious) conditions, essentially lying about the patient's condition. Hence, a patient may also reject high-priced treatments as he knows that the doctor may be lying about his medical condition and overcharging him for a minor problem. Both the patient-dumping practice and the cheating behavior worry policy makers, hurt consumers, and cause significant social welfare loss. Ironically, the doctors may also lose some customers and therefore make less profit than they could have. This motivates our research questions: What impact does the presence of the ethical doctor have on the self-interested doctor? What pricing strategies will different types of doctors adopt at equilibrium? How does the level of ethics in the market (i.e., the probability of the doctor being ethical) affect pricing and market efficiency? Does a more ethical marketplace necessarily lead to higher market efficiency? We study the strategic interactions between the two types of doctors and the patients who are ex ante uncertain about both the severity of their

conditions and about whether the doctor they face is ethical or self-interested. Though we use the healthcare service market as the expositional context, our analyses are applicable to other service industry settings with key characteristics that we have discussed. We provide novel insights into the pricing strategies and market outcomes for services markets such as consulting, investment management, or healthcare services.

Our research contributes to the behavioral economics literature that studies agents with social preferences, i.e., agents who care about factors other than their own profits (e.g., social outcome, other agents' welfare, or perceived fairness). Cui et al. (2007) incorporate the concern of fairness in a conventional channel to study how fairness may affect channel coordination. Ho and Zhang (2008) employ reference-dependent utility specification to study how the presentation of pricing contracts impacts channel outcome. Amaldoss and Jain (2005, 2008) study the impacts of consumers' context dependent or social preferences on strategic firm behavior. Ho and Su (2009) consider peer-induced fairness when agents engage in social comparison.[4] We contribute to this new, active area of research by examining substantively a services market in which the service provider may be interested only in its own profit or may ethically care about the client's wellbeing.[5] We study the impact of the uncertainty about the service providers' social preferences and the implications on pricing strategies and market outcomes.

We also contribute to the literature on asymmetric information. Most signaling games focus on the pure-strategy separating equilibrium (e.g., Balachander, 2001; Desai, 2000; Desai and Srinivasan, 1995; Moorthy and Srinivasan, 1995; Shin, 2005; Simester, 1995; Soberman, 2003) with

---

[4] Ho et al. (2006a, 2006b) and Goldfarb et al. (2012) review some recent developments of behavioral economic modeling involving reference point, loss aversion and social preference and call upon careful applications in managerial decision making with behavioral regularity.

[5] In the aftermath of the 2008 financial crisis, business ethics is considered critically important and has become a heated topic in both policy discussion and academic research ("A Crisis of Ethic Proportions," *Wall Street Journal*, April 21, 2009, http://online.wsj.com/article/SB124027114694536997.html).

the rare exception of Jiang et al. (2011). In contrast to Jiang et al. (2011) where the high-type seller wants to imitate the low-type seller and the pooling outcome is the main focus, in our scenario, the self-interested doctor has an incentive to imitate the ethical doctor because of the patient's increased willingness to pay for a doctor with social preferences. Both pooling and separating equilibria are important to rationalize the market phenomena.[6] Interestingly, we find that at equilibrium the patient may play mixed strategies out of concern that the doctor may have incentives to overcharge. From a theoretical perspective, we study *two-dimensional* private information—the doctor has private knowledge of her own type (ethical or self-interested) and further, upon seeing the patient, she also has private information about the severity of the patient's condition. This is different from Jiang et al. (2011), where the third-party seller has one-dimensional private information about his demand and he has also an unobservable service level (to the platform owner)—hidden action—which creates a moral hazard issue. In our model, we do not have any hidden action—the pricing decision and whether to offer treatment upon seeing a patient are both observed by the patient—but both dimensions of information (about the doctor's ethical type and about the patient's severity condition) are the doctor's private knowledge. Further, we model social utilities—one player may care about another player's wellbeing in addition to her own profit. Our approach of modeling both social preferences and the key characteristics of multi-dimensional private information in a services market such as healthcare or investment management has not been taken by any of the aforementioned literature.

We highlight several key findings from our analysis. First, in a less ethical market (where the probability of the doctor being ethical is low), a separating equilibrium materializes—the self-interested doctor posts a differential price menu whereas an ethical doctor posts the same price for

<hr>

[6] This is also related to the literature on counter-signaling (Feltovich et al., 2002; Mayzlin and Shin, 2011).

both conditions.[7] Realizing the possibility of being overcharged for a minor condition, the patient will occasionally reject the high-price offer from the doctor posting differential pricing.[8] Thus, in equilibrium, in the case of a self-interested doctor, patients of both severity conditions may go untreated with some probability.

Second, in a more ethical market (where the probability of the doctor being ethical is high), in equilibrium, the ethical doctor still adopts a uniform pricing strategy (posting and charging only one price for both severity conditions), and the self-interested doctor will mimic the ethical doctor's pricing strategy.[9] In this pooling scenario, all patients will accept a treatment offer at the pooling-equilibrium price. However, the self-interested doctor will dump patients with a severe condition because ex post it is not profitable for her to treat such patients at the pooling-equilibrium price.[10]

Third, interestingly, the market efficiency, the ratio of the total patient-welfare-loss prevented net of the doctor's costs over the socially-optimal total patient-welfare-loss prevented net of the doctor's costs, may be lower in a more ethical market than in a less ethical market. The

[7] "Separating" in this situation refers to the fact that the doctor's type is revealed by her pricing decision. The patient's true condition, however, might not be fully revealed, because an ethical doctor always treats the patient and a self-interested doctor may play a mixed strategy (lying, with some probability, to a patient with a minor condition that he has a severe condition and must pay the higher price).

[8] In the U.S. healthcare industry (unlike in asset management and consulting service industries), a good health insurance plan may cover a large portion of an insured patient's bill. However, the price of healthcare services plays an important role in patients' decisions about whether to seek treatment, given that tens of millions of Americans are uninsured or underinsured (e.g., http://health.usnews.com/health-news/family-health/articles/2008/06/10/25-million-americans-are-underinsured), a situation even more common in many undeveloped or developing countries.

[9] "Pooling" in this situation refers to the fact that the patient cannot ex ante tell the doctor's type because both types of doctors post the same, single price. Note that, ex post, patients who received treatment still do not know their true conditions, but a patient who was dumped will know at that point that the doctor is self-interested and that he has a severe condition. Further, even if the dumped patient shares his ex post knowledge about the doctor's type through word-of-mouth with future patients, those patients will still go to see the self-interested doctor, because they may have a minor condition which can benefit from treatment. If they have a severe condition and are dumped by the doctor, they do not really lose anything. With information sharing between early patients and later patients, an interesting scenario may arise if we allow for dynamic pricing. After word-of-mouth reveals the doctor's type to later patients, both doctors may adopt differential pricing in the future period since the self-interested doctor can no longer pool with the ethical doctor. Such a dynamic-pricing model with consumer information sharing is beyond the scope of our paper. For such a model in a quality differentiated market, see Jiang and Yang (2012)

[10] In practice, dumping a patient may take several forms: the doctor may claim to be fully booked, the doctor may refer the patient to another specialist or clinic (in a different market or area, for example), or she may use other means to discourage the patient from getting treatment from her.

intuition lies in the fact that a higher ethical level (i.e., a higher probability of the doctoring being ethical) gives the self-interested doctor more incentive to mimic the ethical doctor's uniform price menu, which may induce the self-interested doctor to switch from differential pricing to uniform pricing (i.e., separating to pooling). Note that in both situations, the patients with minor conditions are always treated. In a less ethical market, the patients with severe conditions may occasionally reject treatments when facing the self-interested doctor's differential pricing. However, in a more ethical market, all these severe-condition patients will be dumped by the self-interested doctor, who adopts a uniform pricing strategy to mimic the ethical doctor. Therefore, the market efficiency may actually be lower in a more ethical market than in a less ethical one. Further, as the level of ethics in a market increases, the patients with minor conditions become worse off and those with severe conditions become better off.

The rest of the paper is organized as follows. In section 2, we develop an analytical framework to model the interaction between the service provider and the client. In section 3, we first examine the case when the service cost is less than the consumer's expected loss from the problem. Then we analyze the high service cost case when the social loss rises, derive both differential pricing and uniform pricing equilibria, and compare the implications of corresponding consumer surplus, profit and social welfare. Section 4 concludes the paper.

2. Model

We model a monopoly services market of the nature we discussed before. For ease of exposition, we will present our model in the healthcare services context though our analysis applies to other contexts such as consulting and asset management services.[11] We will refer to the monopolist service

---

[11] A monopoly model in the services setting that we study is in fact quite reasonable. Many services markets are monopolies or monopolistically competitive markets. For example, healthcare services providers, in rural areas, are typically monopolies; typically there is only one clinic or hospital, or one dental services office in each local rural area. In metropolitan areas, there may be multiple providers but they provide specialized services that are potentially

8

provider as the doctor and the consumer or client as the patient. Without loss of generality, we normalize the total number of patients (N) to one. The patient has an uncertain medical condition: with probability β, he has a severe condition, and with probability $1-\beta$, he has a minor condition. A patient with a severe condition—an *H*-type patient—requires a high treatment or service cost C_H from the doctor. A patient with a minor condition—an *L*-type patient—requires a low treatment or service cost $C_L < C_H$ from the doctor. Upon seeing the patient, the doctor will learn the patient's condition at no cost (i.e., the diagnostic cost is normalized to zero), but this information about the patient's condition is the doctor's private information. The patient does not know ex ante his condition's severity, although he knows the prior probabilities of his condition being severe and being minor.[12] If he goes untreated, the patient of type *i* will incur a welfare loss (in the future), denoted by W_i, $i \in \{H, L\}$ with $W_L < W_H$. So if untreated, the patient's ex ante expected welfare loss is $E(W) = \beta W_H + (1-\beta)W_L$. If the patient receives treatment/service from the doctor, his welfare loss will be completely prevented verifiably after the doctor incurs the corresponding service cost C_i. That is, if the doctor treats the patient, the patient will know after treatment that his health has been restored but he may not know the doctor's actual cost.

There are two types of doctors—ethical (type *e*) or self-interested (type *s*). The self-interested doctor maximizes her own monetary profit as in standard economic models, and her utility from treating patient *i* at price *p* is $p - C_i$. In this paper, we will use the terms "utility," "profit," and "payoff" interchangeably. The ethical doctor cares about both profit and patients'

---

differentiated in location, style, and technical competence and they compete mostly monopolistically (Pauly and Satterthwaite, 1981).

[12] In practice, there may be a range of severity conditions. If the condition is severe enough, a patient will probably know that he has a severe condition, because he may see or feel severe symptoms. Our paper, however, focuses on the cases in which the severity of the patient's condition is not yet so drastically serious that the patient can "self-diagnose" the true severity of his condition.

wellbeing, so her utility from treating patient i at price p is $p - C_i + \alpha W_i$, where a positive constant α denotes the ethical doctor's degree of social preference.[13] The ethical doctor's utility from treating a patient increases as α increases.[14] A doctor of either type derives zero utility if she does not treat the patient. It is common knowledge that the doctor is ethical with probability $\gamma \in (0,1)$ and self-interested with probability $1 - \gamma$. The doctor knows her own type whereas the patient knows only the prior probability of each type of doctor. Thus, essentially, the doctor possesses two-dimensional private information—about her own type (whether ethical or self-interested) and about the patient's type (whether the patient's condition is severe or minor).[15] We focus on the case where the two types of doctors are sufficiently different such that the following condition holds:

$$\alpha \geq \max\{\frac{C_H}{W_H}, \frac{C_L}{W_L}, \frac{C_H - C_L}{W_H - W_L}\} \tag{C1}$$

This implies that unlike the self-interested doctor, the ethical doctor has a strong enough social preference that she is willing to treat a patient for free rather than leave a patient untreated. Further, for nontrivial analysis, we assume $W_i > C_i$, $i \in \{H, L\}$, i.e., it is socially efficient to cure both conditions.

---

[13] An alternative for the ethical doctor's utility or payoff function is $p - C_i + \alpha(W_i - p)$, i.e., the ethical doctor also cares about not only whether a patient's health is restored but also how much the patient pays to receive treatment (i.e., the patient surplus). The utility function we have constructed here is in the spirit of Farley (1986), in which physicians are concerned about their patients' wellbeing (i.e., whether their health is restored).

[14] Note that a self-interested doctor corresponds to $\alpha = 0$. Effectively, we have used the terminology of "self-interested" and "ethical" to represent the two types of doctors with different degrees of social preferences. Our model results remain qualitatively the same as long as there is a large enough difference between the degrees of social preferences of the two types of doctors.

[15] Instead of explicitly using the social preference setting, we can adopt a framework in which the doctor is either efficient (with a lower treatment cost) or inefficient (with a higher treatment cost) for any type of patient. Essentially, we absorb the social utility into the doctor's treatment cost and use the cost type to represent the doctor's type rather than using differences in the degree of social preference. So, in the alternative framework, social preferences represent one factor that can lead to differences in the doctor's effective cost efficiency. Such an alternative framework and our social preference framework are conceptually equivalent.

More formally, the game proceeds as follows. Nature determines the patient's type and the doctor's type. With probability β, the patient's medical condition is severe, i.e., he is of the high-cost type (H). With probability $1-\beta$, the patient has a minor condition, i.e., he is of a low-cost type (L). With probability γ, the doctor is ethical (e), and with probability $1-\gamma$ the doctor is self-interested (s). The doctor learns with certainty her true type $j \in \{e, s\}$ whereas the patient knows only the prior probability distribution of the doctor's type. The doctor then posts the price menu, p_{ji}, $i \in \{H, L\}$ with $p_{jH} \geq p_{jL}$, where p_{ji} will be charged for a patient of type i as announced by the doctor. The patient observes the price menu posted by the doctor, and subsequently sees the doctor. Upon seeing the patient, the doctor costlessly learns the patient's true type—whether his condition is severe (H) or minor (L)—and then she offers to treat the patient at some price p from her posted menu, or dumps/rejects the patient.[16] Based on the offered price p and the price menu $\{p_{jL}, p_{jH}\}$, the patient updates his beliefs about his health condition and about the doctor's type, and decides whether to accept the treatment at price p. If he accepts, he pays p to the doctor, who incurs the treatment cost corresponding to the patient's true condition, and the patient's welfare loss is prevented (i.e., his health is restored). The game ends with either treatment or rejection by the doctor or the patient.

The doctor's strategy consists of the price menu $\{p_{jL}, p_{jH}\}$ and the offer that specifies the probabilities that the doctor demands p_{jL}, p_{jH}, or dumps the patient, conditional on the patient's condition (since the doctor observes the patient's condition). Note that different equilibria may yield

---

[16] Note that the doctor learns the patient's type *after* the patient observes the doctor's price menu. This assumption essentially means that the doctor cannot post different price menus to different patients in the case of N>1 patients. (We normalized the number of patients to one.) Our model setting is therefore a much more reasonable framework than the alternative of assuming that the doctor learns a patient's type and then posts a price menu (and subsequently chooses a price from the menu to offer the patient).

the same market outcome. For example, in one equilibrium, the doctor lists two different prices p_{jL} and p_{jH} for different severity conditions but always prescribes p_{jH}. This equilibrium outcome is equivalent to that of another equilibrium in which the doctor lists a single price p_{jH} for both conditions and always prescribes it. To simplify our analysis and exposition, we restrict the doctor to posting only prices that she prescribes with a positive probability. The patient's strategy is a mapping from a doctor's treatment offer to an accept/reject decision about treatment. We analyze Perfect Bayesian Equilibria, which consist of each type of doctor's strategy and payoff, the patient's strategy and payoff, and the patient's beliefs about the doctor's type and about his severity condition.

3. Analysis

For the sake of completeness, we analyze, in Section 3.1, the straightforward case of $E(W) \geq C_H$, i.e., the expected loss is greater than or equal to the treatment cost for the severe condition. Subsequently, we examine our focal case of $E(W) < C_H$, where the self-interested doctor cannot credibly commit to always treating a patient at a price $E(W)$ because she will be better off rejecting the *H*-type patient after learning his condition.

3.1. Aligned Ex Ante and Ex Post Incentive

We first examine the case of $E(W) \geq C_H$. Under this scenario, the doctor's ex ante and ex post incentives to treat a patient at her requested price are aligned. That is, if ex ante the doctor commits to treating a patient of either condition at a specified price, then ex post she has no incentive to reject/dump any patient. We show that the doctor's maximum utility is achieved by posting and charging one price $E(W)$, which leads to all patients' accepting the offer and receiving treatment. This result is easily proved by contradiction. Suppose that there exists another equilibrium price menu or price offer that the doctor can make which will give her a higher utility. Then it must be that the doctor (regardless of her type) collects from the patient an expected total revenue higher

than $E(W)$, which is the amount achieved by posting and charging one price $E(W)$. This implies that at least one type of patient is paying a price higher than $E(W)$ at equilibrium. This means that the equilibrium must be a separating equilibrium, because for a pooling equilibrium, no patient will be willing to pay more than $E(W)$. But for a separating equilibrium, regardless of whether the patient's condition is revealed, the maximum expected revenue that the doctor can collect is still equal to $E(W)$, not higher. Thus, we have shown that the doctor's maximum utility is achieved by posting and charging one price $E(W)$, which leads to the efficient outcome of all patients accepting the offer and receiving treatment.

3.2. Misaligned Ex Ante and Ex Post Incentive

We now examine our focal case of $E(W) < C_H$, which implies $C_L < W_L < E(W) < C_H < W_H$. Under this condition, the self-interested doctor can no longer credibly commit herself to always treating the patient at the price of $E(W)$ because ex post (after learning the patient's type) the self-interested doctor is strictly better off dumping the H-type patient. In contrast, because of her social preferences, an ethical doctor will derive an overall positive utility from treating both types of patients at $E(W)$ and hence will offer such a price and treat both types of patients if patients know that she is the ethical type. Lemma 1 shows that if the doctor's type is common knowledge, then the ethical doctor can still extract the maximum social surplus, achieving the socially optimal result of having all patients treated. However, the self-interested doctor, because of the lack of credible commitment, will make a lower profit and fail to achieve the socially optimal result.

LEMMA 1. *When the doctor's type is common knowledge, the ethical doctor's optimal strategy is to post and charge*

$p_e^* = E(W)$ *for both types of patients whereas the self-interested doctor will post a price menu*

$\{p_{sL}^* = W_L, \ p_{sH}^* = W_H\}$*, then always charge W_H to an H-type patient, and charge W_L to an L-type patient. The*

patient accepts the offer of $p_{sL}^* = W_L$ and $p_e^* = E(W)$ with probability 1, and the offer of $p_{sH}^* = W_H$ with

probability $\delta^* = \dfrac{W_L - C_L}{W_H - C_L}$. *Further, the self-interested doctor earns a lower profit than the ethical doctor (even*

excluding the social utility component).[17]

When the doctor's type is not ex ante known to the patient, one issue arises. The ethical doctor can no longer achieve the maximal utility by posting and charging a uniform price of $E(W)$. At that price, the patient will not accept treatment because a self-interested doctor will also offer him that price only if he has a minor condition. Thus, the patient can infer that, conditional on being offered treatment at a price of $E(W)$, the probability of his condition being severe must be less than the prior probability β, which implies that he should reject treatment since $E(W \mid p = E(W)) < E(W)$. It turns out that when the doctor's type is observed, the self-interested doctor's most profitable uniform-pricing strategy is to charge W_L. However, as Lemma 1 shows, the self-interested doctor will achieve the highest profit by using a differential price menu rather than a uniform price. Interestingly, in this special case of the doctor's type being common knowledge, the self-interested doctor does not cheat (by overcharging patients with minor conditions) at equilibrium.[18]

LEMMA 2. *The ethical doctor is better off offering uniform pricing with some price $p \geq W_L$ compared with offering any differential pricing menu $\{p_L, p_H\}$ with $p_L < p_H$ regardless of whether the patient believes that she is ethical.*

Lemma 2 shows that the ethical doctor prefers uniform pricing to differential pricing regardless of whether the patient believes that she is ethical or not. That is, even if the patient believes (falsely) that the doctor is self-interested, the ethical doctor still prefers offering uniform

---

[17] All proofs not incorporated in the main body of the paper are relegated to the Appendix.
[18] For more discussion on the no-cheating result, refer to Fong (2005), which focuses on identifying the conditions under which a self-interested doctor will cheat.

pricing to offering a differential price menu. This is a result of the ethical doctor's high opportunity cost for losing an L-type patient when he rejects a high-price offer with a positive probability.

Now we analyze the focal case in which the doctor's type is not observed by the patient. We first examine the separating equilibria, in which the two types of doctors post different price menus, from which the patient can determine the doctor's true type. We then solve for the pooling equilibria, in which both types of doctors post the same price menu.

3.2.1. Separating Equilibria (Differential Pricing)

We start with Proposition 1 to characterize all separating equilibria where a differential pricing strategy is adopted by the self-interested doctor and a uniform pricing strategy by the ethical doctor.

PROPOSITION 1. *At separating equilibria, the two types of doctors post different price menus, from which the patient can infer the doctor's true type.*

(a) The self-interested doctor posts the price menu $\{p_{sL}^*, p_{sH}^*\}$, where $p_{sL}^* = W_L$ and $p_{sH}^* = W_H$. She always offers price p_{sH}^* to an H-type patient and price p_{sL}^* to an L-type patient.

(b) The ethical doctor posts a single price p_e^* for both severity conditions, where $p_e^* \in [W_L, \hat{p}_e]$ with

$$\hat{p}_e = W_L + \frac{\beta(W_H - C_H)(W_L - C_L)}{(1-\beta)(W_H - C_L)} ; \text{ she offers treatment at } p_e^* \text{ to both types of patients.}$$

(c) The patient always accepts treatments at offered prices of p_{sL}^* or p_e^* from the price menus above, but accepts p_{sH}^*

only with probability $\delta^* = \dfrac{W_L - C_L}{W_H - C_L}$.

The usual strict belief system about the doctor's type (i.e., any deviation to a uniform menu different from p_e^* comes from a self-interested doctor) supports the separating equilibria. We observe that at a separating equilibrium, the patient is able to infer the doctor's true type from the price menu she posts. The ethical doctor posts a single price p_e^* whereas the self-interested doctor

posts a differential price menu $\{W_L, W_H\}$. Note that, to the self-interested doctor, the separating

outcome is the same as depicted in Lemma 1. To the ethical doctor, however, the posted separating

price is less than her first-best price when her type is known, i.e., $p_e^* < E(W)$. This is because she

has to ensure that the uniform price is low enough that the self-interested doctor will not deviate to

that uniform price. Such a uniform price p_e^* convinces the patient that she is an ethical doctor; the

doctor's promise to treat all patients at p_e^* is therefore a credible commitment and the patient will

accept such an offer with probability 1. Interestingly, at equilibrium, the self-interested doctor has no

incentives to charge the high price to an L-type patient because upon receiving a high price offer,

the patient will randomize his acceptance such that the doctor is indifferent between offering the

high price and the low price to an L-type patient.

Proposition 1 provides a continuum of perfect Bayesian equilibria depending on the patient's

off-equilibrium beliefs about the doctor's type when a uniform price menu is posted. We show here

that any separating equilibrium with $p_e^* < \hat{p}_e$ can be eliminated by the intuitive criterion (Cho and

Kreps, 1987). Suppose that the doctor deviates from p_e^* to some $p_e \in (p_e^*, \hat{p}_e)$. This posted

uniform price is still equilibrium-dominated for the self-interested doctor regardless of what the

patient believes about the doctor's type.[19] Therefore, the patient should not believe that the doctor

who voluntarily made such a deviation can be of the self-interested type with any positive probability.

That is, the strict off-equilibrium belief corresponding to any $p_e^* < \hat{p}_e$ fails the intuitive criterion. It

is easy to verify $p_e^* = \hat{p}_e$ is the only separating outcome that survives the intuitive criterion (the off-

equilibrium belief is that any deviation to a higher uniform price than $\hat{p}_e$ can come from either type

---

[19] In the proof of Proposition 1, we show that the self-interested doctor does not want to deviate to any uniform price
less than $\hat{p}_e$.

of doctors. Clearly, among all separating equilibria the one with $p_e^* = \hat{p}_e$ also gives the ethical doctor

the highest payoff. We will focus on this unique separating outcome for the rest of the paper. We

summarize the analysis as the following.

PROPOSITION 2. *The unique separating equilibrium that survives the intuitive criteria (out of all the equilibria in*

Proposition 1) is also the ethical doctor's most profitable one among all separating equilibria and corresponds to

$p_e^* = \hat{p}_e$.

In a separating equilibrium, the self-interested doctor's profit is the following:

$$\pi_s^{sep}(p_{sH}^*, p_{sL}^*) = \beta \frac{W_L - C_L}{W_H - C_L}(W_H - C_H) + (1 - \beta)(W_L - C_L). \tag{1}$$

From Proposition 2, the ethical doctor sets $p_e^* = \hat{p}_e$ and her payoff is

$$\pi_e^{sep}(p_e^*) = W_L + \frac{\beta(W_H - C_H)(W_L - C_L)}{(1 - \beta)(W_H - C_L)} + [\beta(\alpha W_H - C_H) + (1 - \beta)(\alpha W_L - C_L)]. \tag{2}$$

Simple algebra shows that the ethical doctor's equilibrium payoff, $\pi_e^{sep}(p_e^*)$, increases in β, the

prior probability of a severe condition. In contrast, the self-interested doctor's payoff decreases in

β.

COROLLARY 1. *Under the separating equilibrium, the self-interested doctor's profit decreases in β whereas the*

ethical doctor's profit increases in β.

Moreover, neither doctor's profit depends on the level of ethics (the probability of being

ethical), γ. Under the separating equilibrium, the ethical doctor accepts both types of patient and

the patient accepts with certainty any treatment offer at a price of $p_{sL}^* = W_L$ or p_e^*. Thus, the

expected welfare loss under the separating equilibrium, denoted by W_{sep}, comes only from the

patient's probabilistic rejection of the self-interested doctor's high-price offer of p_{sH}^{*}. W_{sep} is easily computed below.

$$W_{sep} = (1-\gamma)\beta(1-\delta^{*})(W_{H}-C_{H}) = (1-\gamma)\beta\frac{W_{H}-W_{L}}{W_{H}-C_{L}}(W_{H}-C_{H}) \qquad (3)$$

Clearly, the socially-optimal total patient-welfare-loss prevented (net of the doctors' cost) is $W_{max} = \beta(W_{H}-C_{H})+(1-\beta)(W_{L}-C_{L})$. We define the market efficiency ε as the ratio of the actual total welfare loss prevented over the maximal welfare loss prevented: $\varepsilon \equiv \dfrac{W_{max}-W_{sep}}{W_{max}}$.

COROLLARY 2. *In the separating equilibrium, the market efficiency increases in the level of ethics in the market (γ) though neither type of the doctor's payoff depends on γ.*

Corollary 2 shows that as expected, the level of ethics in the market does not affect the doctor's payoff but it influences the market efficiency. In particular, the market efficiency increases as the probability of a doctor being ethical increases.

3.2.2. Pooling Equilibria (Uniform Pricing)

We now examine pooling equilibria and characterize the uniform-pricing equilibria in Proposition 3.

PROPOSITION 3. *Under a pooling equilibrium, both types of doctors post the same price* $p^{*} \in [\hat{p}_{e}, \hat{p}]$ *for different patient conditions where* $\hat{p} = \dfrac{\beta\gamma W_{H} + (1-\beta)W_{L}}{\beta\gamma + (1-\beta)}$.

(a) The self-interested doctor accepts the L-type patient at price p^{*} but dumps the H-type patient.

(b) The ethical doctor accepts any patient at price p^{*}.

(c) The patient always accepts a treatment offer at p^{*} from the doctor's uniform price menu.

These uniform-pricing pooling equilibria are supported by different off-equilibrium beliefs. Below is the belief system that supports all equilibria (both pooling and separating) and that survives

the refinement by the intuitive criterion: The doctor is self-interested if she posts any differential menu or a uniform menu with a price $p > p^*$. The doctor is ethical if she posts a uniform menu with $p \le \hat{p}_e$. If the doctor's uniform price $p \in (\hat{p}_e, p^*]$, the posterior distribution of the doctor type is the same as the prior. If the doctor's uniform price is $p > p^*$, then she is self-interested.

Note that conditional on the patient accepting the treatment offer, both types of doctors have the same incentives to maximize their monetary profits. Therefore, different from Spence's (1973) job market signaling model, the single-crossing property does not hold in our setting and as we show below, the intuitive criterion cannot eliminate the pooling equilibria from Proposition 3. Note that only the ethical type of doctor has the incentive to prove her identity from deviation. From Lemma 2 and Proposition 1, we obtain that the ethical doctor has no profitable deviation to any differential price menu even if the patient believes she is the ethical type. In addition, if the ethical doctor deviates to any price $p < p^*$, she will also make a strictly lower profit than her equilibrium profit even if doing so convinces the patient that she is ethical. This is because at the current price p^*, all patients already accept her treatment offer since

$$E(W \mid p^*) = \frac{\beta \gamma W_H + (1-\beta) W_L}{\beta \gamma + (1-\beta)} \ge p^*$$ and a lower price will not increase the number of patients that

she treats. Lastly, any deviation $p > p^*$ by the ethical doctor is also not a deviation that is equilibrium-dominated for the self-interested doctor under all possible briefs. Thus, our pooling outcomes survive the intuitive criterion.

If the doctor serving the market is always ethical ($\gamma = 1$), she will charge a price of $E(W)$ and will always treat all patients. This equilibrium is socially efficient and allows the doctor to extract the entire social surplus. When the self-interested doctor exists, she has an incentive to mimic the ethical doctor in order to profit from the L-type patient but dump the H-type (high-cost) patient

19

since $p^* < E(W) < C_H$. The self-interested doctor's dumping H-type patients results in a net welfare loss and hence socially suboptimal market inefficiency. In contrast to the extant literature on credence goods (Dulleck and Kerschbamer, 2006; Dulleck, Kerschbamer, and Sutter, 2011), which has focused on the welfare loss from a patient's inability to seek high-price treatments, we show that the frequently-observed patient-dumping practice can be another important source of social welfare loss.

COROLLARY 3. *All the pooling equilibria are equally efficient and the doctor's most profitable equilibrium corresponds to that with* $p^* = \hat{p}$.

As Corollary 3 indicates, the most profitable pooling equilibrium for the doctor is when the belief corresponds to $p^* = \hat{p}$. The doctor's profit increases in price since at pooling equilibria the patient always accepts treatment. If we focus on the most profitable equilibrium, we see that the patient's expected loss from no treatment (conditional on being offered price p^*) increases as the probability of the doctor being ethical increases. Consequently, as γ increases, the patient's willingness-to-pay for treatment (under uniform pricing) increases and the self-interested doctor's incentive to mimic the ethical doctor also increases. Though the ethical doctor treats any patient at the equilibrium price, the self-interested doctor always dumps the H-type patients ex post, so all pooling outcomes have the same market efficiency. Recall that it is socially efficient to have both types of patients treated; the market inefficiency in the pooling equilibria comes from the welfare loss of the self-interested doctor's patient-dumping. In a (uniform-pricing) pooling equilibrium, the L-type patient is always treated whereas the H-type patient will be untreated with probability $1 - \gamma$ (dumped by the self-interested doctor). The welfare loss and the corresponding market efficiency are therefore $W_{pool} = (1 - \gamma)\beta(W_H - C_H)$ and $\varepsilon = 1 - W_{pool}/W_{max}$, respectively, which are

the same across all pooling equilibria even though the division of the social surplus between the patient and the doctor may vary.

3.3. Separating versus Pooling

We now examine how the ethical level (the probability of the doctor being ethical) affects which type of equilibria will be realized and analyze the relative market efficiency across markets with different levels of ethics.

3.3.1. The Role of the Level of Ethics

Let us first examine the doctor's profit under the most profitable pooling equilibrium, where both types of doctor post a uniform price $p^* = \hat{p}$.[20] Since the self-interested doctor accepts only the L-type patient and dumps any H-type patient, her profit is

$$\pi_s^{pool} = (1-\beta)(p^* - C_L) = (1-\beta)\frac{\beta\gamma(W_H - C_L) + (1-\beta)(W_L - C_L)}{\beta\gamma + (1-\beta)}. \tag{4}$$

The self-interested doctor's pooling profit increases in the level of ethics (γ), because the patient's willing to pay increases with γ.

In contrast, the ethical doctor treats both H-type and L-type patients, and her payoff is

$$\pi_e^{pool} = \frac{\beta\gamma W_H + (1-\beta)W_L}{\beta\gamma + (1-\beta)} - [\beta(C_H - \alpha W_H) + (1-\beta)(C_L - \alpha W_L)]. \tag{5}$$

Interestingly, different from the case of the separating equilibrium where the self-interested doctor's profit decreases in β, the self-interested doctor's pooling profit is not monotone in β. Corollary 4 below contrasts with Corollary 1.

COROLLARY 4. *Under the pooling equilibrium, the self-interested doctor's profit increases first then decreases with β, while the ethical doctor's profit still increases with β.*

---

[20] The intuitive reason why the most profitable pooling outcome is most reasonable is that the doctor makes the first move and can credibly induce the most favorable belief system; this argument is well articulated in Jiang et al. (2011).

In the unique separating equilibrium, the patient can identify the doctor's type from her posted price menu. From Proposition 2, the self-interested doctor's profit in the separating equilibria is (1), which does not depend on γ. The ethical doctor's separating payoff at price $p_e^* = \hat{p}_e$ is given by (2). Comparing the doctor's profit under differential pricing with that under uniform pricing for both types of doctors ((5) vs. (2); (4) vs. (1)), we find that the doctor's profit is greater under the pooling outcome than under the separating outcome for both types of doctor when the same condition (6) below holds.

$$\gamma > \gamma^* = \left[\frac{(W_H - C_L)(W_H - W_L)}{(W_L - C_L)(W_H - C_H)} - \frac{\beta}{1-\beta} \right]^{-1} \tag{6}$$

Note that $\dfrac{(W_H - C_L)(W_H - W_L)}{(W_L - C_L)(W_H - C_H)} > 1$ and that if the odd ratio $\dfrac{\beta}{1-\beta} < 1$, i.e., there are more H-type patients than L-type patients, then condition (6) is guaranteed. Proposition 4 immediately follows.

PROPOSITION 4. *There exists γ^* such that a separating equilibrium is manifested when $\gamma < \gamma^*$, and a pooling outcome is manifested when $\gamma > \gamma^*$.*

Even though technically there exist multiple equilibria given the parameter γ, the doctor as the monopolist can exact the highest profit by moving first in the game and posting the pricing menu that gives the highest profit. The comparison of pooling equilibrium (uniform pricing) and separating equilibrium (differential pricing) shows that the uniform pricing strategy may benefit the doctor compared with the differential pricing. Our analysis shows that states like Maryland which impose (regulated) uniform pricing may not hurt the healthcare providers as long as the healthcare industry has a high level of ethics. However, if the ethical level is low, the uniform pricing strategy may be detrimental.

3.3.2. Market Efficiency and Patient Welfare

Now we examine the welfare loss and the market inefficiency in markets with different levels of ethics. As discussed earlier, it is socially efficient to have both types of patients treated since the doctor's cost is smaller than the patient's welfare loss if going untreated.

Recall from Corollary 3 that all pooling equilibria have the same market efficiency. Under the (uniform-pricing) pooling equilibrium, the self-interested doctor will dump the *H*-type patients ex post and the ethical doctor will treat both types of patients. Note that the most profitable pooling equilibrium generates the least welfare loss at the level of $W_{pool} = (1-\gamma)\beta(W_H - C_H)$. Under the separating equilibrium, the self-interested doctor is willing to treat both types of patients with a differential price menu. However, patients occasionally reject the high-price treatment out of concerns about the doctor's misreporting incentive. The ethical doctor still posts and offers a uniform price and treats both types of patients. The welfare loss at the separating equilibrium is computed as $W_{sep} = (1-\gamma)\beta\dfrac{W_H - W_L}{W_H - C_L}(W_H - C_H)$. Figure 1 illustrates how the level of ethics affects the market efficiency. As noted, market efficiency is measured by the ratio of the total patient welfare loss prevented over the maximal welfare-loss prevented (net of the doctor's costs). Intuitively, one may expect that the higher the level of ethics in the market, the higher the market efficiency. However, we find that this is not necessarily the case. [21]

---

[21] The socially-optimal (maximal) welfare-loss prevented is not a function of the level of ethics; therefore the relationship of market efficiency and the market welfare-loss is similar to the normalization by the socially-optimal (maximal) welfare-loss.

Figure 1 Effect of γ on market efficiency (ε)
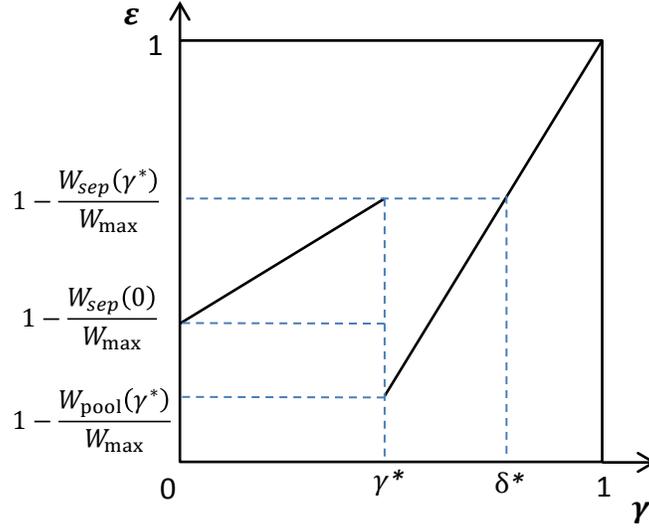
As expected, the welfare loss decreases in the level of ethics γ when $\gamma < \gamma^*$; once the level

of ethics reaches beyond γ^*, the market switches to a pooling regime and the welfare loss again

decreases in γ. That is, under each parameter region, the market efficiency increases. Interestingly,

there is an efficiency gap at $\gamma = \gamma^*$ when the market switches from the differential-pricing

(separating) equilibrium to the uniform-pricing (pooling) equilibrium. This is because under uniform

pricing, the self-interested doctor dumps the patient with severe conditions whereas under

differential pricing, some patients with severe conditions will be treated by the self-interested doctor.

Proposition 5 formally proves that the market efficiency may be lower in a more ethical market than

it is in a less ethical market.

PROPOSITION 5. *Market efficiency may be lower in a more ethical market than in a less ethical one.*

The intuition lies in the fact that a higher ethical level (i.e., a higher probability of the doctor

being ethical) gives the self-interested doctor more incentive to mimic the ethical doctor's uniform

price menu, which may induce the self-interested doctor to switch from differential pricing to

uniform pricing (i.e., separating to pooling). Note that in both situations, the *L*-type patients are

always treated, so the market inefficiency is due to the *H*-type patients. In a less ethical market, *H*-type patients may occasionally reject treatments when facing the self-interested doctor's differential pricing. However, in a more ethical market, *all H*-type patients will be dumped by the self-interested doctor, who adopts a uniform pricing strategy to mimic the ethical doctor. Therefore, the market efficiency may actually be lower in a more ethical market than in a less ethical one.

Next we examine how patients are affected by the level of ethics in the market. Note that the *L*-type patient has zero surpluses in a market with $\gamma < \gamma^*$ since in equilibrium he is offered and accepts the treatment at a price of W_L by the self-interested doctor. And the ethical doctor also offers the *L*-type patient a lower price p_e^* in a less ethical market. But in a more ethical market (with $\gamma > \gamma^*$), the L-type patient *always* pays an equilibrium price $p^* > p_e^* \geq W_L$. Therefore, the *L*-type patient is worse off in the more ethical market. Further, the *H*-type patient obtains an ex ante expected surplus of $\gamma(W_H - p_e^*)$ in the less ethical market with $\gamma < \gamma^*$, but in the more ethical market with $\gamma > \gamma^*$, he obtains an ex ante expected surplus of $\gamma(W_H - p^*)$. One can easily show that in both parameter ranges, the *H*-type patient is better off as γ increases. We summarize the result in the proposition below.

PROPOSITION 6. *The L-type patient is worse off in the more ethical market with $\gamma > \gamma^*$ than in a less ethical one with $\gamma < \gamma^*$ whereas the H-type patient is better off.*

4. Conclusion

In this paper, we have examined the economic and social implications of the particular characteristics of many services markets such as consulting and healthcare. In such markets, the service provider may have more information about the client's problem than the client, and different clients' problems may impose different levels of cost on the service provider. In principle, the

service provider should ethically care about the client's welfare (e.g., by fiduciary duty or ethical codes of conduct), but it is possible that a provider may maximize only its own profit. For example, a doctor may not always heed the welfare of her patients. Some health services providers prescribe more than adequate or necessary levels of care to patients; some even stop providing procedures or services that carry a high liability risk or require a great amount of time and effort. Some simply dump those patients with severe conditions that cost them more to service and focus on the most profitable patients. Furthermore, the client may not know ex ante whether the service provider is ethical or purely self-interested. Using healthcare as the motivating context, we have introduced a game-theoretic model to investigate pricing strategies and the market outcome in such a services market where the service provider has two-dimensional private information—about her own type (whether ethical or self-interested) and about the patient's condition (whether severe or minor).

Our analysis shows several key findings. First, in a less ethical market, a unique separating equilibrium survives the intuitive criterion, and at that equilibrium the self-interested doctor adopts differential pricing whereas an ethical doctor will post the same price for both conditions. Realizing the possibility of being overcharged for a minor condition, the patient will occasionally reject the treatment offer when asked to pay a high price by a doctor using differential pricing. Thus, in the case of a self-interested doctor, both a patient with a severe condition and a patient with a minor condition might go untreated with some probability.

Second, in a more ethical market, the ethical doctor still posts and charges one price for both patient conditions but the self-interested doctor will mimic the ethical doctor's pricing strategy achieving a pooling equilibrium. Under the pooling equilibrium, no patients of either severity condition will reject treatments at the equilibrium price. However, the self-interested doctor will dump patients with a severe condition because it is not profitable for her to treat such patients at the pooling-equilibrium price.

Third, interestingly, the market efficiency may be lower in a more ethical market than in a less ethical one. This is because a higher ethical level gives the self-interested doctor more incentive to mimic the ethical doctor's uniform price menu, which may induce the self-interested doctor to switch from differential pricing to uniform pricing. Since the L-type patients are always treated, the market inefficiency is due to the H-type patients. In a less ethical market, H-type patients may occasionally reject treatments when facing the self-interested doctor's differential pricing. However, in a more ethical market, all H-type patients will be dumped by the self-interested doctor, who adopts a uniform pricing strategy to mimic the ethical doctor. Thus, the market efficiency may actually be lower in a more ethical market. We find that as the level of ethics in a market increases, the L-type patients become worse off and the H-type patients become better off.

To the best of our knowledge, our paper is the first to explore both social preferences and the two-dimensional informational asymmetry in the services setting that we have characterized. Our current study offers several avenues for future research. First, our monopoly analysis is relevant for services markets such as primary care, health services, or consulting services in rural regions, or other monopolistically competitive markets. However, extending our framework to a competitive setting taking into account direct and strategic market competition between service providers may yield additional insights albeit the analysis will become even more complex. Second, it may also be fruitful to examine information sharing between early patients and later patients. An interesting scenario may arise if we allow for dynamic pricing. After word-of-mouth reveals the doctor's type to later patients, both doctors may adopt differential pricing in the future period since the self-interested doctor can no longer pool with the ethical doctor. Third, it may be interesting to investigate the patient's or client's search behaviors or incentives to acquire information in a competitive services market with the same characteristics that we have studied.

References

Amaldoss, W., S. Jain. 2005. Conspicuous consumption and sophisticated thinking. *Management Science* 51(10) 1449-1466.

Amaldoss, W., S. Jain. 2008. Trading up: A strategic analysis of reference group effects. *Marketing Science* 27(5) 932-942.

Ansell, D., R. Schiff. 1987. Patient dumping. *Journal of the American Medical Association* 257(11) 1500-1502.

Balachander, S. 2001. Warranty signalling and reputation. *Management Science* 47(9) 1282-1289.

Cho, I. K., D. M. Kreps. 1987. Signaling games and stable equilibria. *Quarterly Journal of Economics* 102(2) 179-222.

Cui, T., J. Raju, J. Zhang. 2007. Fairness and channel coordination. *Management Science* 53(8) 1303-1314.

Desai, P. S. 2000. Multiple messages to retain retailers: Signaling new product demand. *Marketing Science* 19(4) 381–389.

Desai, P. S., K. Srinivasan. 1995. Demand signaling under unobservable effort in franchising: Linear and non-linear price contracts. *Management Science* 41(10) 1608-1623.

Dulleck, U., R. Kerschbamer. 2006. On doctors, mechanics, and computer specialists: The economics of credence goods. *Journal of Economic Literature* 44(1) 5-42.

Dulleck, U., R. Kerschbamer, M. Sutter. 2011. The economics of credence goods: An experiment on the role of liability, verifiability, reputation, and competition. *American Economic Review* 101(2) 526-555.

Farley, P. J. 1986. Theories of the price and quantity of physician services: A synthesis and critique. *Journal of Health Economics* 5(4) 315-333.

Feltovich, N., R. Harbaugh, T. To. 2002. Too cool for school? Signalling and countersignalling. *RAND Journal of Economics* 33(4) 630-649.

Fong, Y. 2005. When do experts cheat and whom do they target? *RAND Journal of Economics* 36(1) 113-130.

Goldfarb, A., T.-H. Ho, W. Amaldoss, A. L. Brown, Y. Chen, T. H. Cui, A. Galasso, T. Hossain, M. Hsu, N. Lim, M. Xiao, B. Yang. 2012. Behavioral models of managerial decision-making. *Marketing Letters.* 23(2) 405-Ho, T.-H., N. Lim, C. F. Camerer. 2006a. Modeling the psychology of consumer and firm behavior with behavioral economics. *Journal of Marketing Research* 43(3) 307-331.

Ho, T.-H., N. Lim, C. F. Camerer. 2006b. How 'psychological' should economic and marketing models be? *Journal of Marketing Research* 43(3) 341-344.

Ho, T.-H., X. Su. 2009. Peer-induced fairness in games. *American Economic Review* 99(5) 2022–2049.

Ho, T.-H., J. Zhang. 2008. Designing pricing contracts for boundedly rational customers: Does the framing of the fixed fee matter? *Management Science* 54(4) 686-700.

De Jaegher, K., M. Jegers. 2000. A model of physician behaviour with demand inducement. *Journal of Health Economics* 19(2) 231-258.

Jiang, B., K. Jerath, K. Srinivasan. 2011. Firm strategies in the "mid tail" of platform-based retailing. *Marketing Science* 30 (5) 757-775.

Jiang, B., B. Yang. 2012. Quality and pricing decisions in a market with consumer information sharing. Working paper.

Kaufman, A. S., O. Fein, J. J. Fins. 2009. Channeling David E. Rogers, MD: The moral imperative for healthcare reform. *Archives of Internal Medicine* 169(13) 1183-1184.

Mayzlin, D., J. Shin. 2011. Uninformative advertising as an invitation to search. *Marketing Science* 30(4) 666-685.

Moorthy, S., K. Srinivasan. 1995. Signaling quality with a money-back guarantee: The rof transaction costs. *Marketing Science* 14(4) 442-466.

Pauly, M. V., M. A. Satterthwaite. 1981. The pricing of primary cares services: A test of the role of consumer information. *The Bell Journal of Economics* 12(2) 488-506.

Shin, J. 2005. The role of selling costs in signaling price image. *Journal of Marketing Research* 42(3) 302-312.

Simester, D. 1995. Signalling price image using advertised prices. *Marketing Science* 14(2) 166-188.

Soberman, D. A. 2003. Simultaneous signaling and screening with warranties. *Journal of Marketing Research* 40(.2) 176-192.

Spence, M. 1973. Job market signaling. *Quarterly Journal of Economics* 87(3) 355-374.

Vanberg, C.. 2008. Why do people keep their promises? An experimental test of two explanations. *Econometrica* 76(6) 1467-1480.

Wolinsky, A. 1993. Competition in a market for informed experts' services. *RAND Journal of Economics* 24(3) 380-398.

Appendix

PROOF OF LEMMA 1. Since the doctor's type is common knowledge, we simply need to discuss the two cases separately. Let us start with the case of the ethical doctor. From assumption (C1), the ethical doctor's utility for treating a patient of type i at price p is $U_{ei} = p - C_i + \alpha W_i \geq p \geq 0$. Thus, the ethical doctor's ex ante and ex post incentives to treat the patient at price p are aligned. Following the same argument in section 3.1, the ethical doctor's optimal strategy is to post and offer $p_e^* = E(W)$ for both conditions, and the patients always accept the offer p_e^*.

Now we consider the case for the self-interested doctor. Suppose that the self-interested doctor posts a single price p for both conditions. Note that if $p > E(W)$, then no patient will accept that price offer because the patient's expected loss with no treatment is $E(W)$ since he does not know his true condition. If $W_L < p \leq E(W)$, the doctor will offer this price (which is lower than C_H) only to an L-type patient. The patient who receives such a price offer knows that he must have a minor condition, which results in a loss of W_L if untreated. Thus, the patient will reject the offer since $p > W_L$. The patient will accept the offer only if $p \leq W_L$. Thus, the self-interested doctor's most profitable single price is W_L and the corresponding profit is $(1 - \beta)(W_L - C_L)$.

Now we consider the self-interested doctor's most profitable differential price menu. The proof for this special case is similar to that in Fong (2005). Let $\{p_{sL}, p_{sH}\}$ be the price menu, where $p_{sL} < p_{sH}$. The lower price on the menu must satisfy $p_{sL} \leq W_L$ since otherwise no patient, with only prior information about his condition, will accept any price offer from the menu (by similar arguments to the above). In addition, clearly, $p_{sH} \leq W_H$ since otherwise even a perfectly-informed H-type patient will reject that price offer. From the self-interested doctor's perspective, she will offer

prices that will at least cover her cost, that is, $p_{sL} \geq C_L$ and $p_{sH} \geq C_H$. In summary, the self-interested doctor's optimal price menu must satisfy: $C_L \leq p_{sL} \leq W_L$ and $C_H \leq p_{sH} \leq W_H$.

Note that if seeing an H-type patient, the self-interested doctor will offer p_{sH} for sure, because she will incur a loss treating an H-type patient at p_{sL}. However, when seeing an L-type patient, the self-interested doctor may also cheat and offer the high price p_{sH}. If the doctor always offers p_{sH} to an L-type patient (as well as to an H-type patient), then the (uninformed) patient will reject such offers for sure since that is higher than his expected loss. Let ρ be the probability that the self-interested doctor will offer p_{sH} conditional on seeing an L-type patient. And let δ be the probability that the patient will accept the offer p_{sH}. At equilibrium (of the subgame following the posting of a differential price menu), the doctor should be indifferent between offering p_{sH} and offering p_{sL} to an L-type patient. Thus, $\delta = \dfrac{p_{sL} - C_L}{p_{sH} - C_L}$. Similarly, the patient should also be indifferent between accepting and rejecting the high-price offer p_{sH}. Thus, $p_{sH} = E(W \mid p_{sH}) = W_H \Pr(H \mid p_{sH}) + W_L \Pr(L \mid p_{sH})$, where "$\mid p_{sH}$" means conditional on being offered the price p_{sH}.

$$\Pr(H \mid p_{sH}) = \frac{\Pr(p_{sH} \mid H)\Pr(H)}{\Pr(p_{sH} \mid H)\Pr(H) + \Pr(p_{sH} \mid L)\Pr(L)} = \frac{\beta}{\beta + \rho(1 - \beta)},$$

$$\Pr(L \mid p_{sH}) = 1 - \Pr(H \mid p_{sH}) = \frac{\rho(1 - \beta)}{\beta + \rho(1 - \beta)}$$

Thus, the doctor should set $p_{sH} = E(W \mid p_{sH}) = \dfrac{\beta W_H}{\beta + \rho(1 - \beta)} + \dfrac{\rho(1 - \beta)W_L}{\beta + \rho(1 - \beta)}$, from which

we easily obtain $\rho = \dfrac{\beta(W_H - p_{sH})}{(1 - \beta)(p_{sH} - W_L)}$. The self-interested doctor's total expected profit is given by

$\pi_s(\{p_{sL}, p_{sH}\}) = \beta\delta(p_{sH} - C_H) + (1-\beta)\rho\delta(p_{sH} - C_L) + (1-\beta)(1-\rho)(p_{sL} - C_L)$. Substituting the expressions for ρ and δ, we can simplify the self-interested doctor's profit to

$$\pi_s(\{p_{sL}, p_{sH}\}) = \frac{\beta(p_{sL} - C_L)(p_{sH} - C_H)}{p_{sH} - C_L} + (1-\beta)(p_{sL} - C_L).$$

Note that this profit function is a strictly increasing function in p_{sL} and p_{sH} when $C_L \leq p_{sL} \leq W_L$

and $C_H \leq p_{sH} \leq W_H$. Thus, in the relevant range, the optimal prices are $p_{sL}^* = W_L$ and $p_{sH}^* = W_H$, and

the maximum profit is $\pi_s^*(\{p_{sL}^*, p_{sH}^*\}) = \frac{\beta(W_L - C_L)(W_H - C_H)}{W_H - C_L} + (1-\beta)(W_L - C_L)$, which is clearly

larger than her maximum profit from uniform pricing that we found earlier. At this equilibrium,

$\rho^* = 0$ and $\delta^* = \frac{W_L - C_L}{W_H - C_L}$. And lastly, this profit is clearly smaller than the ethical doctor's profit

(even excluding the social-preference component),

$\pi_e^*(p_e^*) = E(W) - [\beta C_H + (1-\beta)C_L] = \beta(W_H - C_H) + (1-\beta)(W_L - C_L)$. □

PROOF OF LEMMA 2. First, let us consider the case in which the patient (mistakenly) believes that

the ethical doctor is a self-interested doctor. Then, when a price menu $\{p_L, p_H\}$ with $p_L < p_H$ is

posted, a patient will accept the high-price offer $p_H \leq W_H$ with probability $\delta = \frac{p_L - C_L}{p_H - C_L}$. Let ρ be

the probability that the ethical doctor offers the high price p_H to an L-type patient, i.e.,

$\rho = \Pr(p_H \mid L)$. Similar to our analysis in Lemma 1, we find

$$\Pr(H \mid p_H) = \frac{\Pr(p_H \mid H)\Pr(H)}{\Pr(p_H \mid H)\Pr(H) + \Pr(p_H \mid L)\Pr(L)} = \frac{\beta}{\beta + \rho(1-\beta)},$$

$$\Pr(L \mid p_H) = 1 - \Pr(H \mid p_H) = \frac{\rho(1-\beta)}{\beta + \rho(1-\beta)}.$$

Thus, to maximize her payoff, the ethical doctor should set

$$p_H = E(W \mid p_H) = \frac{\beta W_H}{\beta + \rho(1-\beta)} + \frac{\rho(1-\beta)W_L}{\beta + \rho(1-\beta)}, \text{ from which we easily obtain } \rho = \frac{\beta(W_H - p_H)}{(1-\beta)(p_H - W_L)}.$$

The ethical doctor's total expected payoff is given by

$$\pi_e(\{p_L, p_H\}) = \beta\delta(p_H - C_H + \alpha W_H) + (1-\beta)\rho\delta(p_H - C_L + \alpha W_L) + (1-\beta)(1-\rho)(p_L - C_L + \alpha W_L).$$

Substituting the expressions for ρ and δ, we can simplify the ethical doctor's payoff function to

$$\pi_e(\{p_L, p_H\}) = \frac{\beta(p_L - C_L)(p_H - C_H + \alpha W_H)}{p_H - C_L} + (1-\beta)(p_L - C_L + \alpha W_L) - \beta\alpha W_L \frac{(W_H - p_H)(p_H - p_L)}{(p_H - W_L)(p_H - C_L)}$$

which is a monotonic increasing function in p_L and p_H. Thus, the ethical doctor's best price menu

to deviate to is $\{p_H = W_H, p_L = W_L\}$. Below we show that the ethical doctor's payoff from uniform

pricing is strictly higher than her payoff under the deviation to the price menu $\{p_H = W_H, p_L = W_L\}$.

Note that the ethical doctor's worst payoff from a uniform pricing strategy is when the patient

mistook her as a self-interested doctor (whether she uses uniform pricing or differential pricing).

That is, regardless of the patient's belief about the doctor being self-interested, there exists some

$p \in [W_L, E(W)]$ such that $\pi_e(p) \geq \pi_e(W_L) = W_L - [\beta(C_H - \alpha W_H) + (1-\beta)(C_L - \alpha W_L)]$. It is easy to

show that $\pi_e(W_L) - \pi_e(\{p_H = W_H, p_L = W_L\}) = \dfrac{\beta(W_H - W_L)(\alpha W_H - C_H + C_L)}{W_H - C_L} > 0$. Thus, the ethical

doctor strictly prefers some uniform pricing strategy over any differential pricing strategy in the case

in which the patient (mistakenly) believes that she is a self-interested doctor.

Now let us consider the second case, in which the patient (correctly) believes that the doctor

is the ethical type. So, when a menu $\{p_L, p_H\}$ with $p_L < p_H$ is posted, a patient will accept the offer

$p_H \leq W_H$ with probability $\delta = \dfrac{p_L - C_L + \alpha W_L}{p_H - C_L + \alpha W_L}$. Hence, the ethical doctor's payoff becomes

$$\pi_e(\{p_L, p_H\}) = \frac{\beta(p_L - C_L + \alpha W_L)(p_H - C_H + \alpha W_H)}{p_H - C_L + \alpha W_L} + (1-\beta)(p_L - C_L + \alpha W_L).$$

Note that $\alpha \geq \dfrac{C_H - C_L}{W_H - W_L}$ by assumption (C1). It is straightforward to show that when

$\alpha \geq \dfrac{C_H - C_L}{W_H - W_L}$, the ethical doctor's payoff above is a decreasing function in p_H, which implies that

the maximum payoff occurs when $p_H \to p_L$. That is, there always exists a uniform pricing strategy

that strictly dominates any differential pricing menu for the ethical doctor. So, the ethical doctor

again will prefer uniform pricing.   □

PROOF OF PROPOSITION 1.   By Lemma 2, the ethical doctor prefers some uniform pricing strategy

$p \geq W_L$ to any differential pricing menu whether the patient believes that she is ethical or not. Thus,

at separating equilibrium, the ethical doctor will choose some uniform price $p_e^* \geq W_L$ and offer it to

any patient, who will accept the treatment offer as long as $p_e^* \leq E(W)$.

By Lemma 1, if the self-interested doctor's type is known to patients, her most profitable

price menu is the differential menu $\{ p_{sL}^* = W_L,\ p_{sH}^* = W_H \}$. At separating equilibrium, the self-

interested doctor's type will be correctly inferred by the patients, and thus at such an equilibrium, the

self-interested doctor should adopt that differential price menu and the patient's acceptance strategy

is as already discussed in Lemma 1. We must show that the self-interested doctor cannot profitably

deviate to the ethical doctor's strategy; that is, $\pi_s^{sep}(p_{sH}^*, p_{sL}^*) \geq \pi_s^{'}(p_e^*)$, where $\pi_s^{'}(p_e^*)$ is the self-

interested doctor's profit when she deviates to the ethical doctor's equilibrium strategy. Note

that $\pi_s^{'}(p_e^*) = (1-\beta)(p_e^* - C_L)$ since the self-interested doctor will dump H-type patients. Further,

$\pi_s^{sep}(p_{sH}^*, p_{sL}^*)$ should be the same as the self-interested doctor's maximum profit computed in the

proof of Lemma 1. Therefore, the condition for non-deviation by the self-interested doctor

becomes $\pi_s^{sep}(p_{sH}^*, p_{sL}^*) = \dfrac{\beta(W_L - C_L)(W_H - C_H)}{W_H - C_L} + (1-\beta)(W_L - C_L) \geq \pi_s'(p_e^*) = (1-\beta)(p_e^* - C_L)$,

which leads to $p_e^* \leq \hat{p}_e \equiv W_L + \dfrac{\beta}{1-\beta}\dfrac{W_L - C_L}{W_H - C_L}(W_H - C_H) < E(W)$.

Note that the last part of the above inequality $\hat{p}_e < E(W)$ is easily proved below.

$W_L + \dfrac{\beta}{1-\beta}\dfrac{W_L - C_L}{W_H - C_L}(W_H - C_H) < \beta W_H + (1-\beta)W_L \Leftrightarrow \beta < 1 - \dfrac{(W_L - C_L)(W_H - C_H)}{(W_H - W_L)(W_H - C_L)} \equiv RHS$. Since

$E(W) < C_H$, i.e., $\beta W_H + (1-\beta)W_L < C_H$, we obtain $\beta < \dfrac{C_H - W_L}{W_H - W_L}$. Further, one can easily verify

algebraically that $\dfrac{C_H - W_L}{W_H - W_L} < RHS$ is always true given our assumption $W_H > C_H > W_L$. Therefore,

$\hat{p}_e < E(W)$.

Clearly, the usual strict belief system about the doctor's type (i.e., any deviation to a uniform menu different from the ethical doctor's equilibrium price p_e^* comes from a self-interested doctor) supports the separating equilibria. Later, we will apply the intuitive criteria to refine the equilibria. □

PROOF OF COROLLARY 1. The ethical doctor's equilibrium profit (2) can be simplified to

$\pi_e^{sep} = W_L + \dfrac{(W_H - C_H)(W_L - C_L)}{(1/\beta - 1)(W_H - C_L)} + (\alpha W_L - C_L) + \beta[\alpha(W_H - W_L) - (C_H - C_L)],$ which clearly

increases in β on $\beta \in (0,1)$. In contrast, the self-interested doctor's profit (1) can be simplified to

$\pi_s^{sep} = \left[1 - \dfrac{C_H - C_L}{W_H - C_L}\beta\right](W_L - C_L)$, which is a decreasing function in β for $\beta \in (0,1)$ since

$0 < \dfrac{C_H - C_L}{W_H - C_L} < 1.$ □

PROOF OF COROLLARY 2. Substituting W_{sep} and W_{max} into the definition of market efficiency, we

obtain $\varepsilon = \dfrac{W_{max} - W_{sep}}{W_{max}} = \dfrac{\left[1 - (1-\gamma)\dfrac{W_H - W_L}{W_H - C_L}\right]\beta(W_H - C_H) + (1-\beta)(W_L - C_L)}{\beta(W_H - C_H) + (1-\beta)(W_L - C_L)}$. Clearly, $\dfrac{\partial \varepsilon}{\partial \gamma} > 0$. $\square$

PROOF OF COROLLARY 4. From (4), we know the self-interested doctor's profit under the pooling

equilibrium is $\pi_s^{pool} = (1-\beta)\dfrac{\beta\gamma(W_H - C_L) + (1-\beta)(W_L - C_L)}{\beta\gamma + (1-\beta)}$.

Therefore $\dfrac{\partial \pi_s^{pool}}{\partial \beta} = \dfrac{(W_H - C_L)\gamma(\beta^2(1-\gamma) - 2\beta + 1)}{(\beta(\gamma-1)+1)^2} - (W_L - C_L)$ and $\dfrac{\partial^2 \pi_s^{pool}}{\partial \beta^2} = -\dfrac{2(W_H - C_L)\gamma^2}{(\beta(\gamma-1)+1)^3}$.

When $\beta \in [0,1]$, $\gamma \in [0,1]$, and $W_H > C_L$, we know $\dfrac{\partial^2 \pi_s^{pool}}{\partial \beta^2} < 0$; so the $\dfrac{\partial \pi_s^{pool}}{\partial \beta}$ is

monotonically decreasing in β. When $\beta = \dfrac{C_H - W_L}{W_H - W_L}$, $\left.\dfrac{\partial \pi_s^{pool}}{\partial \beta}\right|_{\beta = \frac{C_H - W_L}{W_H - W_L}} < 0$. When $\beta = 0$,

$\left.\dfrac{\partial \pi_s^{pool}}{\partial \beta}\right|_{\beta = 0} = \gamma(W_H - C_L) - (W_L - C_L)$. Therefore the self-interested doctor's profit under the pooling

equilibrium increases first then decreases with β.

Now we turn to the ethical doctor. From (6), the ethical doctor's payoff under the pooling

equilibrium is
$$\begin{aligned}\pi_e^{pool} &= \dfrac{\beta\gamma W_H + (1-\beta)W_L}{\beta\gamma + (1-\beta)} - [\beta(C_H - \alpha W_H) + (1-\beta)(C_L - \alpha W_L)]\\ &= \dfrac{\gamma W_H + (1/\beta - 1)W_L}{\gamma + (1/\beta - 1)} + (\alpha W_L - C_L) + \beta[\alpha(W_H - W_L) - (C_H - C_L)]\end{aligned}$$.

We know $W_H > C_H > W_L > C_L$ and $\alpha(W_H - W_L) - (C_H - C_L) \geq 0$, clearly, $\dfrac{\partial \pi_e^{pool}}{\partial \beta} > 0$; therefore the

ethical doctor's payoff increases in β. $\square$

PROOF OF PROPOSITION 3. We need to show that neither type of doctor has a profitable deviation

from equilibrium conditional on the patient's belief system about the doctor's type. The doctor is

self-interested if she posts any differential menu or a uniform menu with a price $p > p^*$. The doctor is ethical if she posts a uniform menu with $p \leq \hat{p}_e$. If the doctor's uniform price $p \in (\hat{p}_e, p^*]$, the posterior distribution of the doctor type is the same as the prior. If the doctor's uniform price is $p > p^*$, then she is self-interested.

Using Bayes' rule, $\Pr(H \mid p^*) = \dfrac{\Pr(p^* \mid H)\Pr(H)}{\Pr(p^* \mid H)\Pr(H) + \Pr(p^* \mid L)\Pr(L)} = \dfrac{\beta\gamma}{\beta\gamma + (1-\beta)}.$

The patient's expected welfare loss if going untreated is thus

$$E[W \mid p^*] = W_H \Pr(H \mid p^*) + (1 - \Pr(H \mid p^*))W_L = \frac{\beta\gamma W_H + (1-\beta)W_L}{\beta\gamma + (1-\beta)}.$$

Therefore, the patient will accept the treatment offer at any uniform price of

$$p^* \in [\hat{p}_e, \frac{\beta\gamma W_H + (1-\beta)W_L}{\beta\gamma + (1-\beta)}].$$

We first show that the self-interested doctor has no profitable deviation from equilibrium conditional on the patient's belief system. The self-interested doctor will make a strictly lower profit than her pooling equilibrium profit if she deviates to a uniform menu with $p < p^*$ since she still treats the same number of the L-type patients but will reduce her profit because $p < p^*$ (she rejects the H-type patient). She also makes a lower profit if she deviates to $p > p^*$, because the patients believe that such deviations are made by a self-interested doctor and hence will reject the offer $p > W_L$. If the self-interested doctor deviates to a differential price menu, her maximum profit is

$$\pi_s'(W_H, W_L) = \frac{\beta(W_L - C_L)(W_H - C_H)}{W_H - C_L} + (1-\beta)(W_L - C_L) < \pi_s^{pool}(p^*) = (1-\beta)(p^* - C_L),$$ as given in

Lemma 1. Note that for all $p > \hat{p}_e$, the pooling profit is larger than the maximum profit from a price menu. Thus, the self-interested doctor will not deviate from p^*.

Now we show that the ethical doctor does not have any profitable deviation. She does not want to deviate to any lower uniform price $p < p^*$ since the patient always accepts her offer p^* anyway. She will make a lower profit (than her pooling equilibrium profit) if she deviates to any higher uniform price $p > p^*$ because she will be believed to be a self-interested doctor and patients will reject treatment at any uniform price $p > W_L$. Similarly, she also makes a lower profit if she deviates to any differential price menu (which leads the patients to believe that she is a self-interested doctor). Thus, we have shown that neither type of doctor will deviate from the equilibrium conditional on the patient's decision based on his belief system. □

PROOF OF PROPOSITION 5. In the proof of Corollary 2, we computed the market efficiency for the separating equilibrium case (where $\gamma < \gamma^*$) and showed $\varepsilon(\gamma)$ is an increasing function of γ. When $\gamma > \gamma^*$, the market is in the (uniform-pricing) pooling equilibrium regime; the welfare loss is given by $W_{pool} = (1-\gamma)\beta(W_H - C_H)$ and $\varepsilon(\gamma) = 1 - \dfrac{W_{pool}}{W_{max}}$ also increases in γ. Note that as illustrated in Figure 1, $\varepsilon(\delta^*) = \lim\limits_{\gamma \nearrow \gamma^*} \varepsilon(\gamma)$, where $\delta^* = \dfrac{W_L - C_L}{W_H - C_L}$ and the limit is taken from the left. Therefore clearly, for any $\gamma \in (\gamma^*, \delta^*)$, there exists some $\gamma^c < \gamma^* < \gamma$ such that for any $\gamma' \in (\gamma^c, \gamma^*)$, $\varepsilon(\gamma) < \varepsilon(\gamma')$. □