DEPARTMENT OF ECONOMICS
JOHANNES KEPLER UNIVERSITY OF LINZ

Price Discrimination via the Choice of Distribution Channels

by

Uwe Dulleck*)
Rudolf Kerschbamer

Working Paper No. 0508
December 2005

Johannes Kepler University of Linz
Department of Economics
Altenberger Strasse 69
A-4040 Linz - Auhof, Austria
www.econ.jku.at

*)uwe.dulleck@jku.at
phone +43 (0)70 2468 -8239, -8238 (fax)

Price Discrimination Via the Choice of Distribution Channels

Uwe Dulleck* and Rudolf Kerschbamer†

This Version: July 2005

Abstract

This article studies the use of different distribution channels as an instrument of price discrimination in credence goods markets. In credence goods markets, where consumers do not know which quality of the good or service they need, price discrimination proceeds along the dimension of quality of advice offered. High quality advice and appropriate treatment is provided to the most profitable market segment only. Less profitable consumers are induced to demand a treatment without a serious diagnosis. If consumers differ in the probabilities of needing different treatments some consumers are potentially overtreated. By contrast, under heterogeneity in the valuations of a successful intervention some consumers are potentially undertreated. Our results help to explain the casual observation that in the early phase of the IT industry only low quality equipment was distributed via warehouse sellers while today it is quite common to see high quality equipment at discounters.

JEL Classifications: L15, D82, D40

Keywords: Price Discrimination, Distribution Channels, Credence Goods, Experts, Discounters

*Department of Economics, University of Vienna, Hohenstaufengasse 9, A - 1010 Vienna, E-Mail: uwe.dulleck@univie.ac.at, Tel.: +43-1-4277 374 27, Fax: +43-1-4277 9 374

†Department of Economics, University of Innsbruck and CEPR, E-Mail: rudolf.kerschbamer@uibk.ac.at, Tel.: +43-512-507 7400, Fax: +43-512-507 2980

1 Introduction

At the dawn of the personal computer age, a serious user typically bought his or her PC from an outlet that helped to choose the right equipment, while the same person would go and purchase a "toy PC" for his or her offspring at a warehouse outlet. As time changed, whereas university departments, firms and many professionals (as architects, attorneys etc.) still buy from outlets that offer some advice in choosing the right quality, no-frills warehouses sell "all-inclusive" high capacity PCs to the mass market. One has the impression that warehouses, not offering any advice or help to choose the efficient PC, switched from providing low quality/capacity equipment to selling high capacity equipment, where high capacity PCs often carry many features never used by a typical consumer.

This paper interprets the market behavior of the IT industry as a manifestation of second degree discriminatory pricing in a market for credence goods. In credence goods markets, where consumers do not know which quality of the good or service they need, price discrimination proceeds along the dimension of quality of advice offered. High quality advice and appropriate treatment is sold to the most profitable market segment only. Less profitable consumers are induced to demand a procedure without a serious diagnosis. If consumers differ in the probability of needing different interventions then low-cost consumers get high quality advice and appropriate treatment while high-cost consumers are potentially *overtreated*; that is, they are induced to demand high quality equipment without a serious diagnosis. By contrast, under heterogeneity in the valuation of a successful intervention high-valuation consumers get efficient diagnosis and appropriate treatment, while low-valuation ones are potentially *undertreated*; that is, they are induced to demand low quality equipment independently of the severity of their problem.

The selling of IT equipment involves credence goods characteristics in the sense that experts can help a customer to determine the cheapest capacity available that enables him to carry out the tasks the customer wants to use the equipment for. Also, in the IT industry there is discrimination in the dimension of quality of advice offered: there are PC manufacturers who distribute their equipment through IT warehouses that offer only selected qualities of equipment at a relatively low price and through specialized dealers who offer the entire assortment as well as advice on choosing the right quality. The explanation offered by the present analysis for the casual observation that in the early phase of the IT industry only low quality equipment was distributed via warehouse sellers while today it is quite common to see high quality equipment at discounters is based upon the dominant heterogeneity among consumers: In the early days of the IT industry the dominant heterogeneity was in the valuation for a successful match while today the dominant heterogeneity is in the quality that fits the intended needs best.

Let us detail our framework and findings. We set up our model in terms of a single expert who offers a menu of tariffs. Consumers observe the menu and decide under which tariff they wish to be served, knowing that the expert's behavior crucially depends on the type of tariff they choose. One menu option the expert might offer is an "expert tariff" meaning that consumers can be sure that if they choose this option they will get serious diagnosis and efficient treatment. The expert might also offer different "non-expert tariffs" meaning that consumers who choose such a tariff can be sure not to get a diagnosis and only a limited array of treatment qualities. Later (in Section 5) we re-interpret our model as one where a single producer uses different distribution channels as a way to price discriminate among consumers. There we show that the situation where a monopolistic experts offers a menu containing "expert" and "non-expert" tariffs can be re-interpreted as a setting where a monopolistic manufacturer chooses different distribution channels – expert outlets and warehouse outlets – for her goods.

The single-expert model is as follows: On the demand side of the market there are many consumers. Each consumer (he) has a problem but does not know how severe it is. He only knows the probability distribution over problems and his valuation for a successful treatment.

On the supply side, there is a single expert (she) – later to be re-interpreted as a single manufacturer.[1] At the outset the expert posts a menu of tariffs. Each tariff specifies a list of repair prices. Consumers observe the menu and then decide whether to visit the expert. If a consumer decides for a visit, he specifies the tariff under which he wishes to be served. Now the expert performs a diagnosis and then recommends a treatment. If the customer ac-

[1] The relevant condition is not that a single expert/manufacturer monopolizes the market, but rather that experts/manufacturers have some degree of market power. In a model in which capacity is required to serve customers (cf. e.g. Emons 1997 and 2001, or Richardson 1999) experts/manufacturers have market power (independently of the number of producers who compete for customers) whenever tight capacity constraints hamper competition. Similarly, consumer loyalty, travel costs together with location, search costs, collusion and many, many other factors might give rise to market power.

cepts the expert provides the recommended treatment and charges the price posted for it.

The expert seller has the opportunity to recommend and provide an inefficient treatment. Two varieties of inefficiency may arise in our model, provision of a high quality repair when a simple treatment would have been sufficient to solve the problem, and provision of a simple procedure when a high quality intervention is needed to fix the problem. We refer to the former kind of inefficiency as over-, to the latter kind as undertreatment. Overtreatment results in a successful intervention while undertreatment does not.[2]

Whether the expert has the right incentive to perform a serious diagnosis and to provide the appropriate treatment depends upon the type of tariff under which a customer is served. If the intervention prices specified by the tariff are such that providing one of the treatments is more profitable than providing any of the others then the expert will recommend and provide the most profitable treatment without a serious diagnosis. Only under equal mark-up tariffs where the differences in the intervention prices reflect the differences in treatment costs will the expert perform a serious diagnosis and recommend the appropriate treatment. Consumers know that; that is, they infer the expert's incentives from the intervention prices.

We study two types of heterogeneity among consumers. First we look at heterogeneity in the expected cost of efficient treatment. We show that without price discrimination the expert may serve less than the efficient number of consumers but whoever is served gets efficient diagnosis and appropriate treatment. If price discrimination is permitted, the number of consumers who get served increases but only low-cost consumers are treated efficiently. High cost ones are potentially *overtreated*; that is, they are induced to buy a high quality intervention without a serious diagnosis. Secondly, we study heterogeneity in the valuation of a successful intervention. In this case high-valuation consumers are treated efficiently while low-valuation ones are potentially *undertreated*; that is, they are induced to demand a low quality intervention without a serious diagnosis.

The intuition for our over- and undertreatment results is as follows: If perfect price discrimination were possible, the expert would provide high quality diagnosis and appropriate treatment to all consumers. This follows

[2]Another kind of problem frequently studied in the literature on credence goods is 'overcharging', that is, charging for a more expensive intervention than provided. Overcharging is ruled out in our model by the assumption that consumers are able to observe and verify the intervention performed by the expert.

from the observation that consumers infer the expert's behavior from treatment prices. So the expert cannot gain by cheating. Consequently the best she can hope for is to appropriate the entire surplus generated by honest behavior. In the absence of information about the identity of consumers perfect discrimination is impossible. With imperfect discrimination the expert sells high quality diagnosis and appropriate treatment at a relative high expected price to the most profitable market segment only. For the rest of the market this policy is unattractive since the expected price is larger than the valuation of a successful intervention. Offering efficient diagnosis and treatment at a lower expected price to the residual demand is impossible because this would induce the more profitable market segment to switch to the cheaper policy. So, the expert offers in addition to the expensive efficient diagnosis and treatment policy a cheaper but also less efficient one. This latter policy is designed in such a way that it attracts the less profitable residual demand while being unattractive for the 'good' segment of the market.

In the setting with heterogeneity in the expected cost, the 'good' segment of the market is the one of low-cost consumers. Thus, this segment gets the efficient diagnosis and treatment policy. And the rest of the market? It is potentially overtreated. For low-cost consumers the overtreatment policy in unattractive since their problem is most likely to be a minor one, implying that buying the expensive efficient diagnosis and provision policy still entails a lower expected cost than buying a major treatment at a bargain price.

Under heterogeneity in the valuation of a successful intervention the 'good' segment of the market is the segment of high-valuation consumers. So, the expert skims off this segment with the efficient diagnosis and provision policy. And the residual demand consisting of low-valuation consumers? It is potentially undertreated. The undertreatment policy is unattractive for high-valuation consumers because they have more to lose if the procedure fails.

Our undertreatment result stands in sharp contrast to the findings of another credence goods paper. In a model in which capacity constrained experts provide procedures to consumers who differ in their valuation of a successful intervention, Richardson (1999) finds that in equilibrium all treated consumers are potentially overtreated.[3] A closer look at his paper reveals

[3] To the best of our knowledge there are only two further contributions with heterogeneous consumers, one of them is the more verbal paper by Darby and Karni (1973), the other the 1993 article by Pitchik and Schotter. In both papers heterogeneity is only used to purify a mixed strategy equilibrium.

that the driving forces behind the Richardson overtreatment and our undertreatment result differ. Our (over- and undertreatment) results are driven by the expert's desire to induce self selection among consumers. By contrast, Richardson's findings result from a lack of power to pre-commit to the price of the major treatment. At the diagnosis stage the expert can tell the consumer that the basic intervention is insufficient to cure his problem and inform him about the additional amount he would have to pay if he accepts an upgrade. The customer knows that the basic intervention might fail. He is therefore prepared to pay some additional amount for a stronger treatment. If this amount exceeds the difference in treatment costs (as it is the case under Richardson's assumptions) then the expert has an incentive to always recommend a stronger treatment.[4]

Other credence goods papers analyze substantially different settings.[5] Emons's (1987 and 1993) work is similar to ours in that it assumes that the type of intervention is verifiable and in that it studies the incentives of experts to under- or overtreat consumers. It is different, however, in that consumers are assumed to be homogeneous. Emons finds that whether the market mechanism induces non-fraudulent behavior depends on the amount of information consumers have at hand to infer the experts' incentives to be honest. Alger and Salanié (2002) study a homogeneous-consumer model in which the degree of verifiability is a continuous variable. They identify an equilibrium in which experts defraud consumers in order to keep them uninformed, as this deters them from seeking a better price elsewhere. Pitchik and Schotter (1987 and 1993), Wolinsky (1993 and 1995) and Taylor (1995) assume that the type of intervention is not verifiable and analyze expert's temptation to overcharge homogeneous customers. Pesendorfer and Wolinsky (2003) investigate a model where effort is needed to diagnose a consumer and where an expert's effort investment is unobservable. Their contribution focuses on the effect an additional diagnosis (by a different expert) has on the consumer's evaluation of a given expert's effort.

Outside the credence goods literature our results have close analogies in the literature on monopolistic screening. In a model in which consumers with different tastes for quality (or service) have unit demands for a good, Mussa and Rosen (1978) show that a monopolist who only knows the aggregate

[4]Here, remember that equal mark-ups are necessary to induce an expert to perform a serious diagnosis and to provide the appropriate treatment.

[5]See Dulleck and Kerschbamer (2003) for a synthesis of some of the findings in this area.

distribution of tastes will in general offer a menu of price quality combinations. As compared to the first best outcome, (i) the monopolist tends to enlarge the range of qualities offered, and (ii) almost all consumers buy lower quality products than would be socially optimal. There are several differences between the Mussa and Rosen model and the setting considered here. First, while there is a natural order in the quality space in the Mussa and Rosen model, there is only a partial order in this dimension in the credence goods setting considered here. In particular, high quality in the Mussa and Rosen model unambiguously corresponds to serious diagnosis and appropriate treatment in the setting considered here. The unusual feature in the case of credence goods is, that there are different (unordered) 'lower quality levels'. As discussed above, depending on whether customers differ in their probabilities to need certain treatments or in their valuation for a successful intervention or both determines whether and which type of low quality service ('no diagnosis and always a major treatment', or 'no diagnosis and always a minor intervention') is used as a screening device. Another difference to the Mussa and Rosen model is that quality (i.e., a serious diagnosis) is not costly in our model.[6] As Acharyya (1998) has shown, Mussa and Rosen's results heavily depend on the cost of quality upgrades. In particular, if quality is not costly, the monopolist will offer only one quality, the best available one, as her optimal policy and the only source of inefficiency that remains is the familiar monopoly pricing distortion. Only for the case of multiple demand Gabszewicz and Wauthy (2002) show that quality discrimination may take place even if provision of quality involves no cost of any sort. As we will see below, consumers can have unit demand and diagnosis can be costless in the case of credence goods, and still the expert may refrain from diagnosing some customer groups.

The next section introduces the basic version of the single-expert model. In the basic version there are only two types of problem and two types of procedures and consumers differ in the expected cost of efficient treatment only. In Section 3 we explore the effects of price discrimination in the basic model. Section 4 discusses several extensions/modifications. First, we extend the basic model to allow for an arbitrary number of problems and arbitrary number of interventions; then we modify our basic model to a setting where

[6] The rationale for assuming zero diagnosis cost is that, if the expert finds it profitable to refrain from providing diagnosis to some consumers when diagnosis costs are zero, then, *a forteriori,* she will do so with positive diagnosis costs.

consumers differ not in the expected cost of efficient treatment but rather in their valuation of a successful intervention; and finally we look at a setting in which consumers differ in both dimensions, in their expected cost of efficient treatment and in their valuation of a successful intervention. In Section 5 we re-interpret the different tariffs offered in the single-expert model as different distribution channels chosen by a monopolistic manufacturer. There, we also revisit some of the modelling assumptions and discuss alternatives. Section 6 concludes.

2 A Basic Model of Credence Goods

On the demand side of the market there is a continuum with mass one of risk-neutral consumers. Each consumer (he) has a problem that needs to be treated. The consumer knows that he has a problem, but does not know how severe it is. He only knows that he has an *ex ante* probability of g_k to have a problem of degree k.

On the supply side of the market there is a single risk-neutral expert (she). Her task is to diagnose consumers, to recommend them a treatment, and, if a consumer accepts, to provide the recommended treatment.

In the basic model of Section 3 there are only two degrees of problem, a minor ($k = 1$) and a major ($k = 2$) one. Each of them can be efficiently treated by exactly one treatment. We denote the type of treatment that efficiently fixes a problem of degree k by c_k. The less severe problem is less costly to be treated. That is, if we denote the cost of the treatment that efficiently fixes a problem of degree k by c_k, then $c_1 < c_2$.[7] The more expensive treatment fixes either problem, while the cheap one is only good for the minor problem.[8]

Table 1 represents the gross utility of a consumer given the type of treatment he needs and the type he gets. If the type of treatment is sufficient, a consumer gets utility $\hat{v}$. Otherwise he gets 0. To motivate this payoff structure consider a car with either a minor problem (car needs oil in the engine) or a major problem (car needs new engine), with the outcomes being 'car

[7] For convenience, both the type of treatment and the associated cost is denoted by c.

[8] In Section 4 we extend the basic model to allow for $n > 2$ degrees of problem ($k \in \{1, .., n\}$). There we assume that, for any $k < l$, problem k is less severe than problem l so that $c_k < c_l$. A more expensive treatment fixes all problems cheaper treatments fix, while the cheapest one is only good for the least severe problem.

Customer's utility		Customer needs	
		c_1	c_2
Customer	c_1	$\hat{v}$	0
gets	c_2	$\hat{v}$	$\hat{v}$

Table 1: Utility from a Credence Good

works' (if appropriately treated or overtreated) and 'car does not work' (if undertreated).[9] The case of overtreatment is the lower left cell (with more than two degrees of problem and more than two treatment qualities the case of overtreatment is the lower left triangle), the case of undertreatment the upper right cell. Note that overtreatment is not detected by the customer ($\hat{v} = \hat{v}$) and hence cannot be ruled out by institutional arrangements. This is not the case with undertreatment; it is detected by the consumer ($\hat{v} > 0$) and might even be verifiable. In the present paper we assume that this is not the case. This means that payments cannot be conditioned on the resolution of the problem. However, consumers are assumed to be able to observe and verify the delivered quality (they know and can prove in which row of the table they are). Thus, payments can be conditioned on the type of treatment provided.

Our consumers are heterogeneous. In the basic model of Section 3 we assume that consumers differ in their probabilities of needing different treatments, but receive the same gross valuation $\hat{v}$ from a successful intervention. More precisely, we assume that each consumer is characterized by his type t and that a consumer of type t has the major problem with probability $g_2^t = t$ and the minor one with probability $g_1^t = 1 - t$. Consumers' types are drawn independently from the same concave c.d.f. $F(\cdot)$, with differentiable strictly positive density $f(\cdot)$ on $[0, 1]$. $F(\cdot)$ is common knowledge, but a consumer's type is the consumer's private information.[10] For further

[9] Of course, not all credence goods have such a simple payoff structure. For instance, the payoff for an appropriately treated major disease might differ from that of an appropriately treated minor disease. Similarly, the payoff for a correctly treated minor disease might differ from that of an overtreated minor disease. Introducing such differences would burden the analysis with additional notation, without changing any of the results, however.

[10] Car owners know how they treat their vehicles and the associated risk of needing certain repairs, auto mechanics know only the distribution. Similarly, patients know their eating and smocking habits and the associated risk of getting certain diseases, doctors only the distribution.

use, we define C^t as type t's expected cost of efficient treatment, i.e., $C^t = \sum_{k=1}^{2} g_k^t c_k = c_1 + t(c_2 - c_1)$.

Each consumer incurs a sunk cost c if he visits the expert independently of whether he is actually treated or not. This cost represents the time and effort cost incurred by the consumer in visiting a doctor, taking the car to a mechanic, etc. Consumers are maximizers of expected utility. The utility to a consumer who has been treated is his gross valuation as depicted in Table 1 minus the price paid for the treatment minus the sunk cost c. For further use, we define v as a consumer's valuation net of sunk cost, i.e., $v = \hat{v} - c$. The utility to a consumer who has not been treated is his reservation payoff, which we normalize to equal zero[11], and the utility to a consumer who has visited the expert but has decided not to receive the recommended treatment is $-c$. By assumption, if a consumer is indifferent between visiting the expert and not visiting the expert, he decides for a visit, and if a consumer who decides for a visit is indifferent between accepting and rejecting the expert's recommendation, he decides for acceptance.

The expert maximizes expected profit. The expert's profit is the sum of revenues minus costs over the customers she treated. By assumption, the expert recommends the appropriate treatment if she is indifferent between recommending the appropriate and recommending the wrong treatment and this fact is common knowledge among all players.[12]

The interaction between consumers and the expert is depicted in Figure 1. This figure shows the game tree for the special case where the monopolistic expert courts a single consumer whose type is known with certainty. The variables $\hat{v}$, c, g_1, g_2, c_1 and c_2 are assumed to be common knowledge. At the outset the expert posts take-it-or-leave-it tariffs. Each tariff specifies the prices p_1 and p_2 for c_1 and c_2, respectively. In the special case covered by the figure the expert posts a single tariff only. With heterogeneous consumers the expert might want to post a menu of tariffs. The consumer observes

[11] Here, the implicit assumption is that the outside option is not to be treated at all. Again, the car example provides a good illustration. A car may be inoperable for a minor or a major reason, with the lack of treatment giving the same outcome ('car does not work') as undertreatment. The medical example behaves differently. For instance, letting a cancerous growth go untreated is much different than letting a benign growth go untreated. See Footnote 9 above, however.

[12] Introducing some guilt disutility associated with providing the wrong treatment would yield the same qualitative results as this common knowledge assumption provided the effect is small enough to not outweigh the pecuniary incentives.

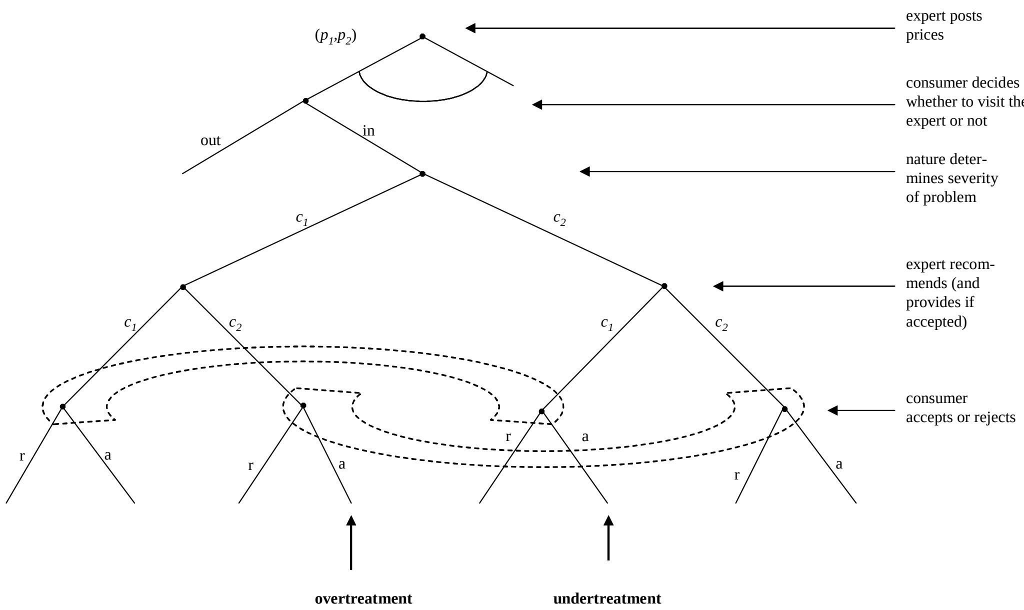

Figure 1: Game Tree for the Basic Model ($t = \bar{t}$ for all consumers, say).

the tariffs and then decides whether to visit the expert or not. If he decides against the visit, he remains untreated yielding a payoff of zero for both players. If he decides for a visit, he specifies the tariff under which he wants to be treated.[13] Then a random move of nature determines the severity of his problem. Now the expert diagnoses the consumer. In the course of her diagnosis she learns the customer's problem and then recommends either the cheap or the expensive treatment. Next the consumer decides whether to accept or reject the recommendation. If he rejects, his payoff is $-c$, while the expert's payoff is zero. If the consumer accepts the expert provides the recommended treatment and charges for it. The game ends with payoffs determined in the obvious way. The extensive form for our model with a continuum of heterogeneous consumers and with a menu of tariffs can be constructed from this game tree in the usual way.

Throughout we restrict attention to situations where the following two conditions hold:

$$v > c_2$$

[13] This move is absent in the figure where the expert posts a single tariff only.

$$c \geq c_2 - c_1$$

The first of these inequalities says that it is efficient to treat both types of problem and the second inequality entails that expert and consumer are in effect tied together once the diagnosis has been made.[14] Relaxing this last restriction complicates the analysis without generating qualitatively different results (provided $c > 0$).[15]

This is the basic setup of our credence goods game. In Section 4 we extend the model to allow for an arbitrary number of problems and an arbitrary number of treatments. There we also analyze the setting where consumers differ in their valuation for a successful intervention.

Throughout the paper we use the following notation: A tariff (p_1, p_2) implies incentives it provides for the expert to perform diagnosis and to provide treatment. Three classes of tariffs are to be distinguished, tariffs that contain a higher mark-up for the expensive treatment $(p_2 - c_2 > p_1 - c_1)$, tariffs that have a higher mark-up for the cheap treatment $(p_2 - c_2 < p_1 - c_1)$, and tariffs with equal mark-ups $(p_2 - c_2 = p_1 - c_1)$. We denote tariffs in the first class by Δ_{02}, tariffs in the second by Δ_{10}, and tariffs in the third by Δ_{12}. As will become clear below, only under tariffs where the differences in the intervention prices reflect the differences in treatment costs (equal mark-up tariffs) will the expert perform a serious diagnosis and recommend the appropriate treatment. Under tariffs where the intervention prices depart from the equal mark-up rule the expert will recommend and provide the most profitable treatment without a serious diagnosis. That is, the expert will perform a honest diagnosis and provide the appropriate treatment under a Δ_{12} tariff, she will always recommend and provide the cheap treatment under a Δ_{10} contract and she will always recommend and provide the expensive treatment under a Δ_{02} contract.[16] For convenience we will often denote not only a specific tariff but also the implied mark-up by Δ. That is, the term Δ will then stand for the mark-up on the treatment that is provided under the respective contract $(\Delta = \max\{p_1 - c_1, p_2 - c_2\})$.

[14] In Subsection 4.1. where there are $n > 2$ degrees of problem the relevant condition is $c \geq c_n - C^0$, in Subsection 4.2 where consumers differ in their gross valuation $\hat{v}$ but have the same g_2 the relevant condition is $c \geq (1 - g_2)(c_2 - c_1)$.

[15] See Section 5 for a discussion.

[16] We use the terms tariff, price-vector and contract interchangeably.

3 Price Discrimination in the Basic Model: Overtreatment

In this section we analyze the effects of price discrimination in our basic model. Before beginning we present a benchmark result for a setting in which the expert cannot price discriminate among consumers. Without price discrimination the expert chooses equal mark-up prices and serves her customers honestly. If the difference in the expected cost between the best and the worst type is small relative to the efficiency gain of treating the worst type then the expert serves all consumers. Otherwise prices are such that some consumers do not consult her even though serving them would be efficient. This is nothing but the familiar monopoly-pricing inefficiency: The monopolistic expert would like to appropriate as much of the net gain from treatment as possible but, because of heterogeneous consumers, she puts up with the risk of losing some consumers in order to get a higher price from the remaining ones. We record the monopoly pricing result in Proposition 1.

Proposition 1 *Consider the basic model with two degrees of problem and consumers who differ in their probabilities of needing different treatments only. Suppose the monopolistic expert cannot price-discriminate among customers. Then, in the unique perfect Bayesian equilibrium (PBE), the expert posts and charges equal mark-up prices* ($p_k - c_k = \Delta$ *for* $k = 1, 2$)*. If the difference in expected cost between the best and the worst type is large relative to the efficiency gain of treating the worst type* [$c_2 - c_1 > (v - c_2)f(1)$] *then prices are such that high-cost consumers decide to remain untreated* ($\Delta > v - C^t$ *for* t *strictly higher than some* $\bar{t} \in (0, 1)$)*, while all other types visit the expert* ($\Delta \leq v - C^t$ *for* $t \leq \bar{t}$) *and get serious diagnosis and appropriate treatment. Otherwise all consumers are efficiently served under equal mark-up prices* ($\Delta \leq v - C^t$ *for all* t).

Proof. Obvious from the discussion below and therefore omitted. ■

The equal mark-up result is readily illustrated graphically. First notice that under verifiability the expert's incentives in performing diagnosis and in providing treatments depend upon the type of tariff under which a customer is served. If the intervention prices specified by the tariff are such that providing one of the treatments is more profitable than providing the other then the expert will recommend and provide the more profitable treatment without a serious diagnosis. So, if we fix the mark-up for the minor intervention at $p_1 - c_1$ and increase the mark-up for the major intervention from 0 (as it is done in Figure 2) then the expert's incentives remain unchanged over the interval $(0, p_1 - c_1)$: she will always recommend and provide the minor treatment at the price $p_1 = c_1 + \Delta_{10}$. Consequently, the expected utility of a consumer of type t is constant in this interval at $v - C^t - t(\hat{v} - c_2 + c_1) - \Delta_{10}$, where the term $t(\hat{v} - c_2 + c_1)$ reflects the efficiency loss from undertreatment. Similarly, if we start at $p_1 - c_1$ and increase $p_2 - c_2$ then the expert will always recommend and provide the major treatment at the price $p_2 = c_2 + \Delta_{02}$. So, the consumer's utility in this segment is $v - C^t - (1-t)(c_2 - c_1) - \Delta_{02}$, where the term $(1-t)(c_2 - c_1)$ reflects the efficiency loss from overtreatment. Only at the single point $p_2 - c_2 = p_1 - c_1 = \Delta_{12}$ where the difference in the intervention prices reflects the difference in treatment costs is the expert indifferent between the two types of treatment and, therefore, behaves honestly.[17] So, at this point there is no efficiency loss and type t's expected utility jumps discontinuously upward to $v - C^t - \Delta_{12}$.

Consumers infer the expert's incentives from the intervention prices. So the expert cannot gain by cheating. Consequently, the best she can do is to post an equal mark-up tariff and to behave honestly. With equal mark-ups the monopolistic expert is interested in two variables only, in the magnitude of the mark-up $\Delta = \Delta_{12}$ and in the number of visiting consumers. The result then follows from the observation that the expert's problem is nothing but the familiar monopoly pricing problem for revenue per customer Δ and demand curve $D(\Delta) = F[(v - c_1 - \Delta)/(c_2 - c_1)]$.[18]

For our next result we allow the expert to post more than one tariff. Since consumers' tastes differ, the monopolist might want to target a specific tariff for each consumer or at least different tariffs for different consumer groups. However, in the absence of information about the identity of each consumer (the expert only knows the aggregate distribution of probabilities

[17] The assumption that it is common knowledge among players that the expert provides the appropriate treatment whenever she is indifferent plays an important role in Proposition 1 in generating a unique PBE. Without this assumption there exist other PBE which are supported by the belief that the expert deliberately mistreats her customers under each equal mark-up vector - or, that the expert deliberately mistreats her customers under equal mark-up prices that are too high. We regard such equilibria as implausible (see Footnote 12 above) and have therefore introduced the common knowledge assumption which acts as a restriction on consumers' beliefs.

[18] That the condition $(c_2 - c_1) > (v - c_2)f(1)$ is necessary and sufficient for an interior solution is easily verified by checking that the expert's profit is an increasing function of Δ at $\Delta = v - c_2$ if and only if this condition is satisfied.

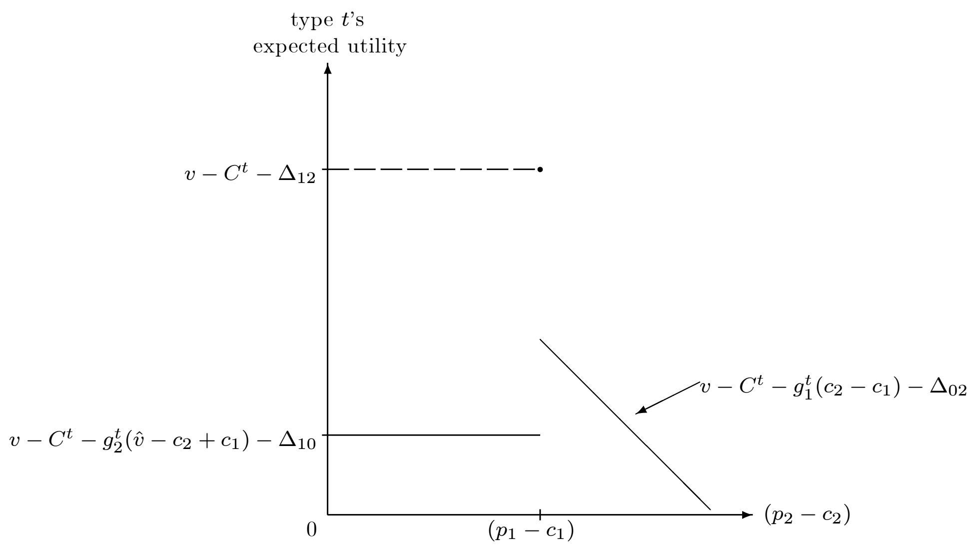

Figure 2: Type t's Expected Utility under Different Price Vectors

of needing different treatments) the expert must make sure that consumers indeed choose the tariff designed for them and not the tariff designed for other consumers. This puts self-selection constraints on the set of tariffs offered by the monopolist. As our next result shows, the monopolist uses the quality of advice offered as a self selection device.

Proposition 2 *Consider the basic model with two degrees of problem and consumers who differ in their probabilities of needing different treatment only. Suppose that the expert can price discriminate among consumers (rather than being restricted to post a single tariff only). Then, in any PBE, the expert posts two tariffs, one with equal mark-ups, and one where the mark up for the major treatment exceeds that for the minor one.*[19] *Both tariffs attract customers and in total all consumers are served. Low cost consumers are served under the former tariff and always get honest diagnosis and appropriate treatment; high-cost consumers are served under the latter and always get the major treatment, sometimes inefficiently.*

Proof. The proof proceeds in four steps. In *Step 1* we first show that any arbitrary menu of tariffs partitions the type-set into (at most) three

[19] The menu may contain some redundant tariffs too, i.e., some tariffs that attract no consumers.

subintervals delimited by cut-off values t_{10}, t_{12} and t_{02} with $0 \leq t_{10} \leq t_{12} \leq t_{02} \leq 1$ and either $t_{12} = t_{02}$ or $t_{02} = 1$ (or both) such that (i) the optimal strategy of types in $[0, t_{10})$ is to choose a Δ_{10} tariff, (ii) the optimal strategy of types in $[t_{10}, t_{12}]$ is to decide for a Δ_{12} tariff, and (iii) the optimal strategy of types in $(t_{12}, 1]$ is either to choose a Δ_{02} tariff ($t_{02} = 1$), or to remain untreated ($t_{12} = t_{02}$).[20] Our strategy is then to show in *Step 2* that an optimal price-discriminating menu cannot have $t_{10} = t_{12}$ (that is, there must be an equal mark-up tariff which attracts a strictly positive measure of types), to show next (in *Step 3*) that $t_{10} = 0$ whenever $t_{10} < t_{12}$ (that is, the expert has never an incentive to post a menu where both an equal mark-up tariff and a tariff with a higher mark-up for the cheap treatment attract types), and to show in the end (*Step 4*) that the expert has indeed always a strict incentive to cover a strictly positive interval by a tariff with a higher mark-up for the expensive treatment ($t_{12} < t_{02} = 1$).

Step 1 First note that any arbitrary menu of tariffs can be represented by (at most) three variables, by the lowest $\Delta_{02} = p_2 - c_2$ from the class of Δ_{02} tariffs (we denote the lowest Δ_{02} in this class by Δ_{02}^l), by the lowest $\Delta_{10} = p_1 - c_1$ from the class of Δ_{10} tariffs (we denote the lowest Δ_{10} in this class by Δ_{10}^l), and by the lowest equal mark-up Δ_{12} from the class of all equal mark-up tariffs in the menu (denoted by Δ_{12}^l).[21] To see this, note that with $n = 2$ each possible price vector is member of exactly one of these three classes, and that a customer who decides for a vector in a given class will always decide for the one with the lowest Δ.[22] An immediate implication is that each menu of tariffs partitions the type-set into the above mentioned three subintervals. This follows from the fact that the expected utility under Δ_{02}^l is type-independent (implying that either $t_{12} = t_{02}$ or $t_{02} = 1$ or both),

[20] The borderline types t_{10} and t_{12} are indifferent between the strategies of the types in the adjacent intervals (whenever such intervals exist). Here note that we allow for $t_{12} = 1$ (all consumers are served and no consumer chooses a Δ_{02} tariff), for $t_{10} = t_{12}$ (no consumer is attracted by a Δ_{12} tariff), and for $t_{10} = 0$ (no consumer is attracted by a Δ_{10} tariff). Price discrimination requires, however, that at least two of the three relations hold as strict inequalities.

[21] An immediate implication of this observation is that successful price discrimination requires that some types are mistreated with strictly positive probability. Why? Since at least two tariffs must attract a positive measure of consumers and since only one of them can be an equal mark-up tariff.

[22] Under a Δ_{02} tariff neither the consumer nor the expert cares about the associated p_1. All tariffs in the group that have the same Δ_{02} can therefore be thought off as being a single tariff without any loss in generality. The argument for Δ_{10} tariffs is symmetric.

while the expected utility under both the Δ_{12}^l tariff and the Δ_{10}^l tariff is strictly decreasing in t, and from $\hat{v} > c_2 - c_1$ (implying that the Δ_{10}^l-function is steeper than the Δ_{12}^l-function).

Step 2 To see that $t_{10} < t_{12}$, suppose to the contrary that $t_{10} = t_{12}$. Then $t_{10} > 0$, since $t_{10} = t_{12} = 0$ is incompatible with price-discrimination (and since - by Proposition 1 - a non-price-discriminating expert will always decide for a Δ_{12} vector). But such a menu is strictly dominated, since the Δ_{10}^l vector can always be replaced by a vector with equal mark-ups of $\Delta_{12} = \Delta_{10}^l + g_2^{t_{10}}(\hat{v} - c_2 + c_1)$; the latter attracts exactly the same types as the replaced one and yields a strictly higher profit.

Step 3 To see that $t_{10} = 0$ whenever $t_{10} < t_{12}$, suppose to the contrary that $0 < t_{10} < t_{12}$. Then the expert's profit is strictly increased by removing all Δ_{10} vectors from the menu. This follows from the observation that (by the monotonicity of the expected utility $-$ in t $-$ under Δ_{12}) all types in $[0, t_{10})$ switch to Δ_{12}^l when all Δ_{10} vectors are removed from the menu, and from the fact that the expected profit per customer is strictly higher under Δ_{12}^l than under Δ_{10}^l whenever $0 < t_{10} < t_{12}$, since $\Delta_{12}^l \leq \Delta_{10}^l$ is incompatible with the shape of expected utilities ($\Delta_{12}^l \leq \Delta_{10}^l$ implies that $v - C^t - \Delta_{12}^l > v - C^t - \Delta_{10}^l - g_2^t(\hat{v} - c_2 + c_1)$ for all $t > 0$ contradicting $t_{10} > 0$). Thus, $t_{10} = 0 < t_{12} \leq t_{02} \leq 1$. So, if price discrimination is observed in equilibrium it is performed via a menu that contains two tariffs, one with equal mark-ups, and one with a higher mark-up for the more expensive treatment.[23]

Step 4 We now show that the expert has always a strict incentive to post such a menu. Consider the equal mark-up vector posted by the expert under the conditions of Proposition 1. The mark-up in this vector is at least $\Delta_{12} = v - c_2$, in an interior solution even higher. First suppose that the monopolist's maximization problem under the conditions of Proposition 1 yields an interior solution (i.e., $\Delta_{12} > v - c_2$). Then the expert can increase her profit by posting a menu consisting of two vectors, the one chosen under the conditions of Proposition 1 and a Δ_{02} vector with $p_2 = v$. The latter vector guarantees each type an expected utility equal to the reservation utility of 0. Thus, all types that remain untreated under the conditions of Proposition 1 will opt for it since they are indifferent. Also, all types served under the conditions of Proposition 1 still choose the equal mark-up vector since $v - C^t$ is strictly decreasing in t. Hence, since $v > c_2$, and since all types

[23] The menue might contain some redundant vectors too, which can safely be ignored, however.

in $[0,1]$ have strictly positive probability, the expert's expected profit is increased.[24] Now suppose that the monopolist's maximization problem under the conditions of Proposition 1 yields the corner solution $\Delta_{12} = v - c_2$. Then again the monopolist can increase her profit by posting a menu consisting of two tariffs, a Δ_{02} contract with $p_2 = v$, and a Δ_{12} contract that maximizes $\pi(\Delta_{12}) = \Delta_{12}F[(v - c_1 - \Delta_{12})/(c_2 - c_1)] + (v - c_2)(1 - F[(v - c_1 - \Delta_{12})/(c_2 - c_1)])$. Since $\pi(\Delta_{12})$ is strictly increasing in Δ_{12} at $\Delta_{12} = v - c_2$ an interior solution is guaranteed. ■

Under the conditions of Proposition 2 the expert posts two tariffs, an equal mark-ups to skim-off low-cost consumers and a less profitable overtreatment tariff to serve the rest. Consumers served under the former tariff get honest diagnosis and appropriate treatment, consumers served under the latter always get the expensive treatment, sometimes inefficiently.

Figure 3 illustrates the result. This figure shows how consumers' expected utility under different tariffs varies in the type. First notice that consumers' expected utility under Δ_{02} (where the expert always provides the expensive intervention) is type-independent, while the expected utility under both the Δ_{12} tariff (where the expert behaves honestly) and the Δ_{10} tariff (where the expert undertreats the customer) is strictly decreasing in t. Next notice that the Δ_{10}-function is strictly steeper than the Δ_{12}-function. This follows from the observation that under Δ_{10} higher types have a higher probability that the intervention fails (leading to a loss of $\hat{v}$) while under Δ_{12} they have only a higher probability to get charged for the more expensive treatment (leading to an additional cost of $c_2 - c_1 < \hat{v}$). Finally remember from the discussion of Figure 2 that if a Δ_{10} tariff, a Δ_{12} tariff and a Δ_{02} tariff simultaneously attract customers (as it is the case in the situation depicted in Figure 3) then the inefficient tariffs Δ_{10} and Δ_{02} must have lower mark-ups than the efficient one. Consequently, a situation as depicted in Figure 3 can never arise in equilibrium: the expert could always remove the Δ_{10} tariff from the menu; then all types in $[0, t_{10})$ would switch to Δ_{12} and the expert's profit would be increased. Also, a menu where only a Δ_{10} tariff and Δ_{02} tariff attract

[24] Here note that the expert can do even better by increasing Δ_{12}^l. This follows from the observation that the expert's trade-off under the conditions of Proposition 1 is between increasing the mark-up charged from the types in the segment of served customers and losing some types to the unprofitable segment of not served consumers, while the trade-off here is between increasing the mark-up charged from the types in the segment of customers served under the more profitable equal mark-up vector and losing some types to the segment of customers served under the less profitable Δ_{02}^l vector.

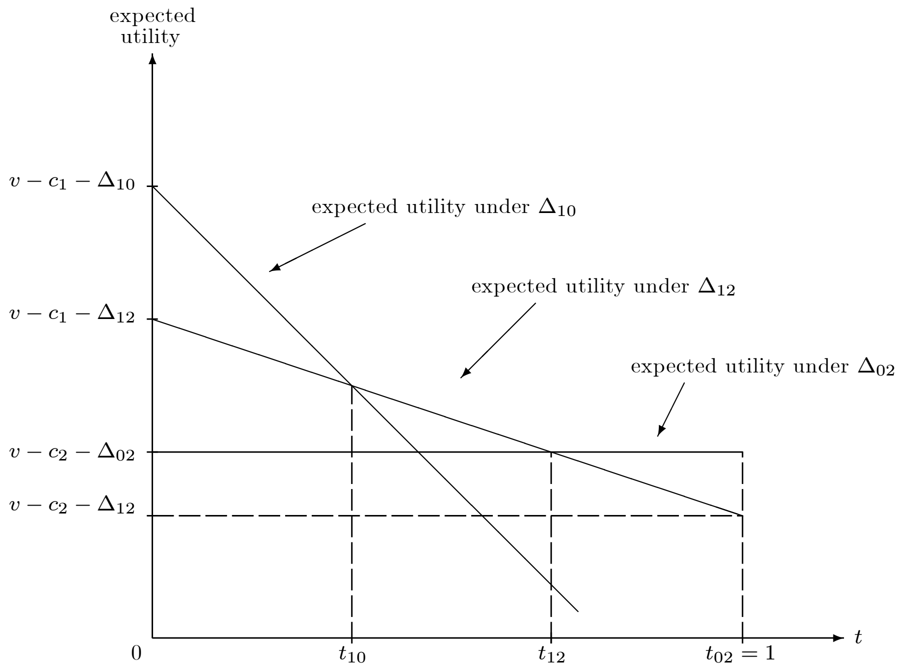

Figure 3: Type Dependent Expected Utilities with $C^t = c_1 + t(c_2 - c_1)$

types can never be an equilibrium menu: the expert could always replace the Δ_{10} tariff by an efficient tariff such that the highest type attracted by Δ_{10} is exactly indifferent between Δ_{10} and Δ_{12}; the Δ_{12} tariff would attract exactly the same types as the replaced Δ_{10} contract and yield a strictly higher profit. So, if price discrimination is observed in equilibrium it is performed via a menu that contains two tariffs, one with equal mark-ups, and one with a higher mark-up for the more expensive treatment.

That the expert has indeed always an incentive to post such a menu is easily seen. First suppose that some consumers remain unserved under the conditions of Proposition 1. Then the expert can increase her profit by posting a menu consisting of two vectors, the one chosen under the conditions of Proposition 1 and a Δ_{02} tariff that leaves zero rents to consumers ($p_2 = v$). Since $v > c_2$, and since all types in $[0, 1]$ have strictly positive probability, the expert's expected profit is increased. Next suppose that the expert's maximization problem under the conditions of Proposition 1 yields the corner solution $\Delta_{12} = v - c_2$. Then again she can increase her profit by posting a menu consisting of two tariffs, a Δ_{02} contract with $\Delta_{02} = v - c_2$, and an efficient tariff with $\Delta_{12} = v - c_2 + \varepsilon$. Still all consumers are served. Those served under Δ_{02} leave the expert with exactly the same profit as before, those served under Δ_{12} are served more profitable. Hence, the expert's profit is again increased.

The equal mark-up in the tariff posted under the conditions of Proposition 2 is strictly higher than that in the tariff of Proposition 1. This follows from the observation that the expert's trade-off is between increasing the mark-up charged from the types in the segment of served customers and losing some types to the unprofitable segment of not served consumers in the latter case, while the trade-off here is between increasing the mark-up charged from the types served under the more profitable equal mark-up vector and losing some types to the segment of customers served under the less profitable second vector. So, some consumers who always get honest diagnosis and appropriate treatment under the conditions of Proposition 1, get (with strictly positive probability) the wrong treatment when the expert can price discriminate among consumers. So, if the difference in expected cost between the best and the worst type is small (so that the monopolist serves all consumers if price discrimination is not permitted) then allowing discrimination unambiguously reduces efficiency. On the other hand, when some consumers are excluded under the conditions of Proposition 1, then there is a trade-off between increasing the number of treated consumers and serving the treated customers efficiently. Overall efficiency might increase or decrease with price discrimination depending on the shape of the distribution function $F(\cdot)$, the valuation v and the cost differential $c_2 - c_1$. As our next result shows, the mass of consumers that are efficiently served under non-discrimination and inefficiently under discrimination increases in the net valuation v and decreases in the cost differential $c_2 - c_1$. At the same time, the mass of consumers that are not served under non-discrimination and served under discrimination decreases in the net valuation and increases in the cost differential. So, price discrimination is *ceteris paribus* more likely to be efficiency enhancing if consumers' valuation of an efficient treatment is small and if the cost differential is large.

Proposition 3 *Consider the basic model with two degrees of problem and with consumers who differ in their probabilities of needing different interventions only. Let* $1 - F(t_{12}^a)$ *stand for the mass of consumers that are not served under non-discrimination and served under discrimination. Similarly, let* $F(t_{12}^a) - F(t_{12}^b)$ *stand for the mass of consumers that are efficiently served under non-discrimination and inefficiently under discrimina-*

tion. Then $1 - F(t^a_{12})$ *increases in* c_2 *and decreases in* v *and in* c_1*, while* $F(t^a_{12})-$ $F(t^b_{12})$ *decreases in* c_2 *and increases in* v *and in* c_1*.*

Proof. The non-discriminating monopolist maximizes $\pi(t_{12}) = F(t_{12})[(v - c_1 - (c_2 - c_1)t_{12}]$ over t_{12}. If this problem has an interior solution t^a_{12} then it satisfies $v - c_1 = (c_2 - c_1)[t^a_{12} + F(t^a_{12})/f(t^a_{12})]$, where $f(.)$ stands for the density function associated with the concave c.d.f. $F(.)$. Therefore, t^a_{12} increases in v and in c_1, and it decreases in c_2. The discriminating monopolist maximizes $\pi(t_{12}) = F(t_{12})[(v - c_1 - (c_2 - c_1)t_{12}] + [1 - F(t^a_{12})](v - c_2) = v - c_2 + (c_2 - c_1)(1 - t_{12})t_{12}F(t_{12})$ over t_{12}. The solution to this problem, denoted by t^b_{12}, depends only on $F(.)$, and not on v, c_2, c_1. Thus, $1 - F(t^a_{12})$ increases in c_2 and decreases in v and in c_1, while $F(t^a_{12})-$ $F(t^b_{12})$ decreases in c_2 and increases in v and in c_1. ■

The following examples illustrates the result:

Example 1 *Suppose the distribution function* $F(.)$ *is given by* $F(x) = x^{1/y}$ *for* $y = 1, 2, ...$ *Then* $\Delta^a_{12} = y(v - c_1)/(y+1)$, $\Delta^b_{12} = [y(v - c_1) + (v - c_2)]/(y+1)$, $t^a_{12} = (v - c_1)/[(y+1)(c_2 - c_1)]$, $t^b_{12} = 1/(y+1)$, *and* $t^a_{12}-$ $t^b_{12} = (v - c_2)/[(y+1)(c_2 - c_1)]$. *So, if* $\hat{v} = 12$, $c_2 = 5$, $c_1 = 2$, $c = 4$ *and* $y = 1$ *(implying that* $F(.)$ *is the uniform distribution) then the non-price-discriminating expert will serve all consumers efficiently under the equal mark-up tariff* $\Delta^a_{1,2} = 3$ $(t^a_{12} = 1)$*. If she is allowed to price-discriminate then she serves half of the population under the equal mark-up vector* $\Delta^b_{1,2} = 4.5$ $(t^b_{12} = 0.5)$*, and the rest under the 'overtreatment tariff'* $\Delta^b_{0,2} = 3$.[25] *So, with this parameter constellation welfare is definitely decreasing when moving from non-discrimination to discrimination because under nondiscrimination all consumers are treated efficiently (*$W^a = 4.5$*) whereas with discrimination customers in the interval* $t \in (0.5, 1]$ *are potentially overtreated, i.e., they receive with probability* $(1-t)$ *an unnecessary expensive treatment (*$W^b = 4.125$*). If* c_2 *increases from* 5 *to* 7 *then the non-price-discriminating expert serves 60% of the consumers efficiently (*$t^a_{12} = 6/10$*) and the rest remains unserved. With this constellation welfare is higher under discrimination (*$W^b = 2.875$*) than under non-discrimination (*$W^a = 2.7$*) because the gain of customers not treated under*

[25] Note that $\Delta^a_{12} = \Delta^b_{02}$ is due to the fact hat g^t_2 has full support on $[0, 1]$ and that parameters are such that all consumers are treated under non-discrimination. Whenever some customers remain untreated under non-discrimination, mark-ups differ. Also, if g^t_2 < 1 for all t then $\Delta^a_{12} \neq \Delta^b_{02}$.

non-discrimination (those in the interval $(0.6, 1]$*) outweights the loss of consumers that are efficiently served under non-discrimination and inefficiently under discrimination (those in the interval* $(0.5, 0.6]$*).*[26] *Similarly, if we start from the same starting point and reduce* v *from* 8 *to* 6 *then the non-price-discriminating expert serves 66% of the consumers efficiently* ($t_{12}^a = 2/3$) *and the rest remains unserved. Again, welfare is higher under discrimination (*$W^b = 2.125$*) than under non-discrimination (*$W^a = 2$*).*

4 Extensions and Modifications

In this section we discuss several extensions/modifications. First, we extend the basic model to allow for an arbitrary number of problems and an arbitrary number of interventions. It turns out that our main result that price discrimination entails potential overtreatment of high-cost consumers extends in this direction. Next we modify our basic model to a setting where consumers differ not in the expected cost of efficient treatment but rather in their valuation for a successful intervention. We show that in this setting the expert provides serious diagnosis and appropriate treatment only to high-valuation consumers while low-valuation ones are potentially undertreated; that is, they are induced to demand a simple procedure without a serious diagnosis. Finally we look at a setting in which consumers differ in both dimensions, in their expected cost of efficient treatment and in their valuation of a successful intervention. It turns out that the expert will always serve at least some consumers efficiently. The rest may get unnecessary or insufficient procedures or no treatment at all.

4.1 More than Two Degrees of Problem: Different Degrees of Overtreatment

In this subsection we extend our analysis to $n > 2$ degrees of problem ($k \in \{1, .., n\}$). We denote the type of procedure that efficiently fixes a problem of degree k by c_k. Without loss of generality we assume that if $k < l$ then problem k is less severe than problem l. Again we assume that a less severe problem is less costly to be treated ($c_k < c_l$ for $k < l$) and that a

[26] Here note that the efficiency gain of treating a type t consumer under Δ_{12} is $v - C^t$, while the efficiency gain of (over-)treating a type t customer under Δ_{02} is $v - C^t - (1 - t)(c_2 - c_1)$.

more expensive treatment fixes all problems cheaper treatments fix, while the cheapest one is only good for the least severe problem. As in the basic model each consumer is characterized by his type t and a type t consumer has probability $g_k^t = g^t(c_k) \geq 0$ of needing treatment c_k, with $\sum_{k=1}^n g_k^t = 1$. Let $G^t(.)$ be the associated c.d.f., i.e., $G^t(c_l) = \sum_{k=1}^l g^t(c_k)$. Also, let C^t denote the associated expected cost of efficient treatment, i.e., $C^t = \sum_{k=1}^n g_k^t c_k$. For the formal analysis we need some structure on the type set. What we want to have is (i) a continuum of types, (ii) for each type t a strictly positive probability of having a problem of degree k $(= 1, ..., n)$, and (iii) an ordering on the type set such that for any two types s and t with $s \leq t$ the probability of having a problem of at least degree k is higher under $G^t(.)$ than under $G^s(.)$ for every degree of problem. A simple way to get such a structure is to take two c.d.f.s $G^1(.)$ and $G^0(.)$ with densities $g^1(.)$ and $g^0(.)$ that have full support on $\{c_1, ..., c_n\}$ such that the former first-order stochastically dominates the latter (i.e., $1 - G^1(.) > 1 - G^0(.)$ for all c_k, or equivalently $G^1(.) < G^0(.)$ for all c_k), and to let the c.d.f. of problem degrees for a type t consumer be given by $G^t(.) = (1-t)G^0(.) + tG^1(.)$. In the sequel we follow this way and assume that consumers' types are drawn independently from the same distribution $F(\cdot)$, with differentiable strictly positive density $f(\cdot)$ on $[0, 1]$. Again, $F(\cdot)$ is assumed to be common knowledge, but a consumer's type is the consumer's private information.

In an $n \geq 2$ framework there are $2^n - 1$ classes of tariffs to consider, the class of equal mark-up tariffs (denoted by $\Delta_{1,2,...,n-1,n}$) and $2^n - 2$ classes of tariffs that have a lower mark-up for at least one and at most $n - 1$ treatments. We denote tariffs that have a lower mark-up for treatment k by $\Delta_{1,..,k-1,0,k+1,...,n}$.[27] For instance, for $n = 3$, a Δ_{103} vector has $p_1 - c_1 = p_3 - c_3 > p_2 - c_2$. Similarly, for $n = 4$, a Δ_{0004} tariff has $p_4 - c_4 > p_k - c_k$ for $k = 1, 2, 3$. The expert's behavior under the n classes of $\Delta_{0,..,0,k,0,..,0}$ tariffs and under $\Delta_{1,2,...,n-1,n}$ is obvious. She will always provide treatment k under tariffs in the former classes, and she will always provide the appropriate treatment under tariffs in the latter class. What about the rest? Our assumption that the expert acts in her customers' interest whenever she is indifferent implies that she uses the cheapest highest mark-up treatment that fixes the problem whenever such a treatment exists. If none of the highest

[27] Note the slight change in the Δ-notation: When we consider an arbitrary number n of problems we insert commas between the different treatments to avoid confusion; that is, we write $\Delta_{1,2,...,n-1,n}$ instead of $\Delta_{12...n-1n}$ and $\Delta_{1,..,k-1,0,k+1,...,n}$ instead of $\Delta_{1...k-10k+1...n}$.

mark-up treatments fixes the problem, then the expert provides the cheapest highest mark-up treatment. For instance, under $\Delta_{0,..,0,k,0,..,0,l,0,..0}$ the expert will provide procedure c_k for problem degrees $h \leq k$, procedure c_l for problem degrees $h \in \{k+1,..,l\}$, and again procedure k for problem degrees $h > l$.

Given these specifications the net utilities of consumers under all possible tariffs are well defined and we can try to extend the arguments for the $n = 2$ to the $n > 2$ case. As is easily verified, Proposition 1 continues to hold if we replace the condition $c_2 - c_1 > (v - c_2)f(1)$ by $C^1 - C^0 > (v - C^1)f(1)$. The result of Proposition 2 generalizes as follows to the $n > 2$ case:

Proposition 4 *Consider the extended basic model with $n \geq 2$ degrees of problem and with consumers who differ in their probabilities of needing different procedures only. Suppose that some consumers remain unserved under the conditions of Proposition 1 (where the expert is restricted to post a single tariff only).*[28] *Then, in any PBE of the game in which price discrimination is permitted, the expert will post a menu in which at least two tariffs attract types, one with equal mark-ups, and at least one tariff with lower mark-ups for cheaper treatments.*[29] *In total all consumers are served. Low cost consumers are served under the former tariff and always get honest diagnosis and appropriate treatment, high-cost consumers are served under (one of) the latter(s) and are never under- but sometimes overtreated.*

Proof. The proof parallels that of Proposition 2 and is available on request. ■

Proposition 4 confirms that our main result that price discrimination entails potential overtreatment of high-cost consumers extends to the setting with $n > 2$ degrees of problem: Again, low-cost consumers are efficiently served under an equal mark-up tariff and the rest of the market gets unnecessary procedures with strictly positive probability. Also again, no kind of undertreatment is observed in equilibrium; that is, under all tariffs offered, each customer will always get an intervention that fixes his problem.

The most important change when moving from the two to the more than two types of problem setting is that there is no longer a guarantee that the

[28] In opposition to the basic model this condition is needed to make sure that price discrimination is observed in equilibrium. The reason is, that in the current setting the boundary solution has $\Delta_{1,2,...,n-1,n} = v - C^1 > v - c_n$, while the boundary solution in the basic model has $\Delta_{12} = v - C^1 = v - c_2$.

[29] Again, the menu may contain some redundant tariffs too, i.e., some tariffs that attract no consumers.

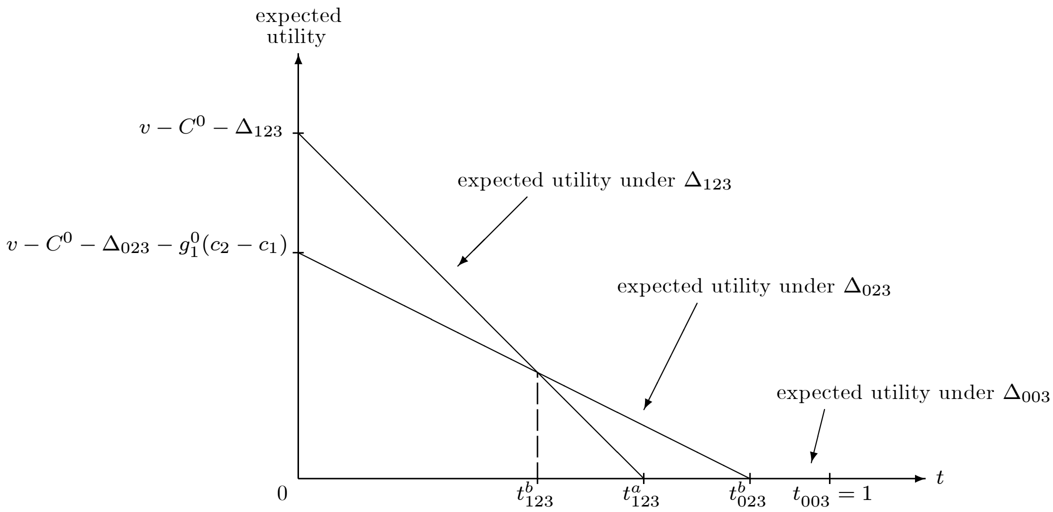

Figure 4: Type Dependent Expected Utilities under $C^t = \sum_{k=1}^{3} g_k^t c_k$

price-discriminating expert will post exactly two tariffs, one with equal mark-ups and one with a higher mark-up for the most advanced intervention. The only guarantee we have is that the expert will post in addition to the equal mark-up tariff at least one other tariff, and that each posted contract other than the equal mark-up tariff will provide the expert with incentives to never under- and to sometimes overtreat customers (i.e., the intervention provided is always sufficient to fix the problem but sometimes a more expensive intervention is provided when a cheaper one would have been sufficient to solve the problem). To get sharper results we would need more information on the shape of the distribution functions and on the cost differential between the different treatments. To see why, look at Figure 4. This figure illustrates the $n = 3$ case. Let us start with a non-discrimination setting in which low-cost consumers (with $t \leq t_{123}^a$) are efficiently served under the equal mark-up tariff Δ_{123} while high-cost consumers (with $t > t_{123}^a$) remain untreated. If we now introduce a Δ_{003} tariff that leaves zero rents to customers ($\Delta_{003} = v - c_3$) then the expert's profit is unambiguously increased. The reason is, that the Δ_{003} contract is flat in the expected-utility/type space; that is, it provides the same utility to all consumers. So all consumers attracted by this contract can be held to their reservation utility. Using a Δ_{023} tariff instead of Δ_{003} has one advantage and one disadvantage. The advantage is, that it is more profitable than the Δ_{003} contract since it entails a smaller inefficiency. The disadvantage is that the tariff is not flat; that is, it offers rents to lower cost consumers. So some consumers (in the figure the market segment $[t_{123}^b, t_{123}^a]$) who would choose the equal mark-up tariff Δ_{123} under the two contract menu $(\Delta_{123}, \Delta_{003})$ will switch to the less profitable Δ_{023} contract if this tariff is also available. (Here notice that if Δ_{123}, Δ_{023} and Δ_{003} attract types, then $\Delta_{123} > \Delta_{023} > \Delta_{003}$.) So, whether it is profitable to post the Δ_{023} tariff in addition to (or instead of) the Δ_{003} contract depends on the magnitude of the two effects, and the magnitude of the two effects depends on the shape of $G^0(.)$, $G^1(.)$ and $F(.)$ and on whether the cost differential $c_{k+1} - c_k$ is increasing or decreasing in k.

4.2 Differences in the Valuation: Undertreatment

Up to now we have investigated settings where consumers differ in their probabilities of needing different treatments only. Now we modify our assumptions and analyze a model where consumers differ in their valuation of a successful intervention $\hat{v}$, but have the same probabilities of needing different procedures. More precisely, we assume that a consumer of type t has valuation $\hat{v}^t = \hat{v} - t$ and that consumers' types are drawn independently from the same concave c.d.f. $F(\cdot)$, with differentiable strictly positive density $f(\cdot)$ on $\left[0, \bar{t}\right]$. Again, $F(\cdot)$ is assumed to be common knowledge, but a consumer's type is the consumer's private information.

With this specification a type t consumer's expected utility under Δ_{12} is $v^t - C - \Delta_{12}$, where $v^t = \hat{v}^t - c$. Similarly, a type t consumer's expected utility under Δ_{02} is $v^t - C - \Delta_{02} - g_1(c_2 - c_1)$. Finally, a type t consumer's expected utility under Δ_{10} is $v^t - C - \Delta_{10} - g_2(\widehat{v}^t - c_2 + c_1)$.

As is easily verified, Proposition 1 continues to hold if we replace the condition $c_2 - c_1 > (v - c_2)f(1)$ by $(v^{\bar{t}} - C)f(\bar{t}) < 1$. Proposition 2 changes to:

Proposition 5 *Consider the basic model with two degrees of problem and two treatment qualities. Suppose that consumers differ in their valuation of a successful intervention $\hat{v}$ (rather than in their probabilities of needing different treatments). Then, if price discrimination is observed in equilibrium, it is performed via a menu containing two tariffs, one with equal mark-ups, and one where the mark up for the minor intervention exceeds that for the major one. High valuation consumers are served under the former tariff and always get serious diagnosis and appropriate treatment; lower valuation consumers* *are served under the latter and always get the minor treatment, sometimes inefficiently.*

Proof. First observe that any arbitrary menu of tariffs partitions the type-set into (at most) three subintervals delimited by cut-off values t_{02}, t_{12} and t_{10} with $0 \leq t_{02}, t_{12} \leq t_{10} \leq 1$ and either $t_{02} = 0$ or $t_{12} = 0$ (or both) such that (i) either the optimal strategy of types in $[0, t_{02})$ is to choose a Δ_{02} tariff (if $t_{02} > 0$), or the optimal strategy of types in $[0, t_{12})$ is to choose a Δ_{12} tariff (if $t_{12} > 0$), (ii) the optimal strategy of types in (t_{12}, t_{10}) is to decide for a Δ_{10} tariff, and (iii) the optimal strategy of types in $(t_{10}, 1]$ is to remain untreated. This follows from the fact that the expected utility under any of these tariffs is strictly decreasing in t, and from the fact that the Δ_{12} and the Δ_{02} function have exactly the same steepness in the expected-utility/type space and that they are both strictly steeper than the Δ_{10} function (see Figure 5 below). The rest of the proof is similar to that of Proposition 2 the only exception being that the Δ_{10} function is not completely flat so that price discrimination may not be observed in equilibrium even if some consumers are excluded under the conditions of Proposition 1 (where the expert is restricted to post a single tariff only). ■

Proposition 5 tells us that in the model where consumers differ in their valuation for a successful intervention but have the same expected cost of efficient treatment, price discrimination entails potential undertreatment of low-valuation consumers.

An explanation is easily provided. Since consumers are homogeneous in the expected cost of efficient treatment an overtreatment tariff, if attractive for low-valuation consumers, will also attract high-valuation ones and hence cannot be used for discriminatory purposes.[30] An undertreatment tariff, on the other hand, is unattractive for high-valuation consumers because they have more to lose if the procedure fails. It is therefore potentially, but not necessarily, useful for discrimination.

Figure 5 provides an explanation for why in the heterogeneity-in-the-valuation-case price discrimination is not necessarily observed in equilibrium even if some consumers are excluded under non-discrimination: If we start with a non-discriminating setting in which the expert posts an equal mark-up

[30]Referring to Figure 5, one can observe that the Δ_{02} and the Δ_{12} tariff have the same steepness in the expected-utility-type space. Given that the mark-up of the Δ_{02} tariff needs to be lower than the mark-up of the Δ_{12} tariff to attract some customers, Δ_{02} tariffs are strictly dominated.

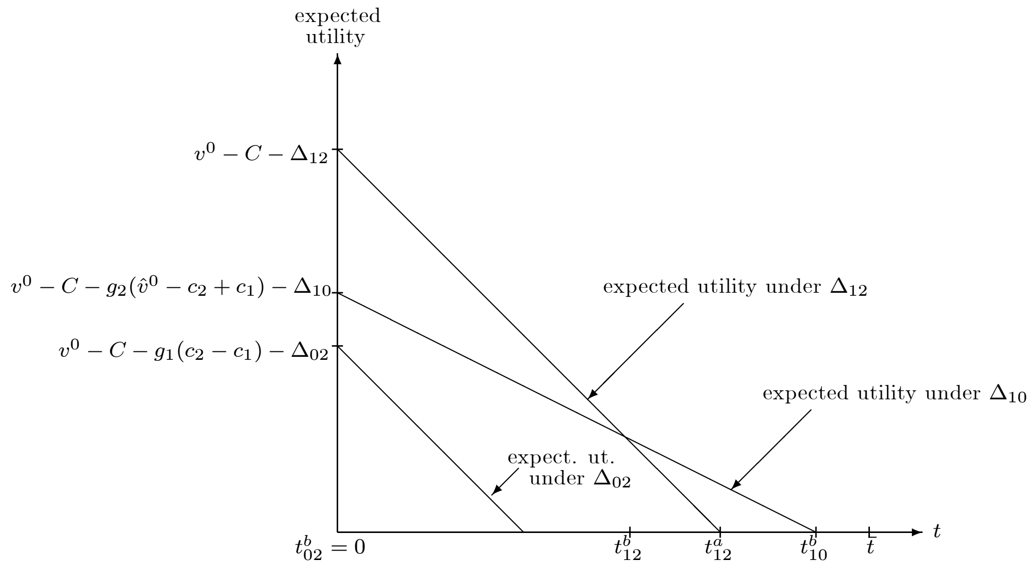

Figure 5: Type Dependent Expected Utilities with $v^t = v - t$

tariff Δ_{12} only, and introduce Δ_{10} as a second tariff then the expert profits because some new consumers (those in the interval $[t_{12}^a, t_{10}^b]$) are attracted. At the same time the expert loses because some consumers (those in the interval $[t_{12}^b, t_{12}^a]$) who used to buy under the more profitable equal mark-up tariff Δ_{12} switch to the less profitable Δ_{10} tariff. Whether the overall effect is positive or negative depends on parameter constellations, that is, on the shape of the distribution function $F(.)$, on the size of the valuation $\hat{v}$ and on the intervention costs c_1 and c_2.

4.3 Two Dimensional Type Set: Over- and Undertreatment

Our previous results suggest that in a setting with a two-dimensional type set over- and undertreatment might coexist in equilibrium. This is indeed the case as the (discrete) example below shows. Before considering this example we first show that even in a two-dimensional world the expert will always treat at least a subset of consumers efficiently.

Proposition 6 *Consider the basic model with two degrees of problem and two treatment qualities. Suppose that consumers differ in their valuation of a successful intervention and in their probabilities of needing different treatments. Further suppose that each consumer has a strictly positive probability* *of having each of the different problems.*[31] *Then, in any PBE the expert will post a menu in which an equal mark-up tariff attracts a non-empty subset of types.*

Proof. Suppose there is no equal mark-up tariff which attracts a strictly positive measure of types. Then, among the tariffs chosen by a strictly positive measure of types, take the one with the highest mark-up for the provided treatment and denote it by Δ_h. Two cases have to be distinguished:

If Δ_h is of the Δ_{02} variety denote the type with the *highest* g_2 among the types attracted by Δ_h by t_h. Then replace Δ_h by a Δ_{12}-tariff such that type t_h is exactly indifferent between Δ_h and Δ_{12}; that is, $\Delta_{12} = \Delta_h + g_1^{t_h}(c_2 - c_1)$. Since consumers with a lower g_2 gain more by the replacement than the critical type t_h, all types attracted by Δ_h under the original menu will be attracted by Δ_{12} under the new menu. Types not attracted by Δ_h under the original menu will either switch to the more profitable Δ_{12} tariff or will choose the same tariff as before the replacement. Thus, since $\Delta_{12} > \Delta_h$ the new menu yields a strictly higher profit.

If Δ_h is of the Δ_{10} variety denote the type with the *lowest* $g_2(\widehat{v} - c_2 + c_1)$ among the types attracted by Δ_h by t_h. Then replace Δ_h by a Δ_{12}-tariff such that type t_h is exactly indifferent between Δ_h and Δ_{12}; that is, $\Delta_{12} = \Delta_h + g_2^{t_h}(\widehat{v}^{t_h} + c - c_2 + c_1)$. Since consumers with a higher $g_2(\widehat{v} - c_2 + c_1)$ gain more by the replacement than the marginal type t_h, all types attracted by Δ_h under the original menu will be attracted by Δ_{12} under the new menu. Types not attracted by Δ_h under the original menu will either switch to the more profitable Δ_{12} tariff or will choose the same tariff as before the replacement. Thus, since $\Delta_{12} > \Delta_h$ the new menu yields a strictly higher profit.[32] ∎

Let us now discuss the example announced earlier. In this example all consumers are efficiently served under equal mark-up prices if the expert can post a single tariff only. With price discrimination the expert uses a Δ_{12}-tariff to skim-off high-valuation/low-cost consumers, a Δ_{10}-tariff to under-

[31] If consumers need the cheap procedure for sure ($g_1^t = 1$) then the tariffs Δ_{10} and Δ_{12} are indistinguishable from an efficiency point of view. Similarly, for consumers who need the expensive treatment for sure ($g_2^t = 1$) Δ_{02} and Δ_{12} are indistinguishable from an efficiency point of view. So, to guarantee that the expert will post a menu in which an equal mark-up tariff attracts a nonempty subset of types at least some consumers must have $g_1^t \in (0, 1)$.

[32] Here notice that the same proof-technique could be used to prove that the result continues to hold if we allow for an arbitrary number of problems and an arbitrary number of interventions.

treat low-valuation/low-cost consumers, and a Δ_{02}-tariff to overtreat high-valuation/high-cost consumers. Low-valuation/high-cost consumers remain unserved with price discrimination although treating them would be efficient.

Example 2 *There are two degrees of problem* ($n = 2$). *Each consumer is characterized by his two-dimensional type* $(g_2^t, \widehat{v}^t)$. *Consumers' types are independently drawn from an equal probability distribution on the discrete support* $\{(0.5, 2.8), (0.2, 2.8), (0.9, 4.0), (0.5, 4.5)\}$. *Each consumer's sunk cost is one* ($c = 1$), *the cost of the expensive treatment is one* ($c_2 = 1$), *and the cost of the cheap treatment is zero* ($c_1 = 0$). *If the expert can post a single tariff only, then she efficiently serves all consumers under the equal mark-up contract* $\Delta_{12} = 1.3$. *With this policy she earns an expected profit of 1.3 per consumer. If the expert can price discriminate among consumers then she increases her expected profit to 1.425 per consumer by posting three price vectors, the equal mark-up tariff* $\Delta_{12} = 2.5$, *the 'overtreatment tariff'* $\Delta_{02} = 2.0$, *and the 'undertreatment tariff'* $\Delta_{10} = 1.2$. $(0.5, 4.5)$*-consumers are served efficiently under the equal mark-up tariff,* $(0.9, 4.0)$*-consumers are potentially overtreated under* Δ_{02}, $(0.2, 2.8)$*-consumers are potentially undertreated under* Δ_{10}, *and* $(0.5, 2.8)$*-consumers remain untreated.*

5 Discussion

In this section we revisit some of the modelling assumptions and discuss alternatives. We also re-interpret the different tariffs offered in the single-expert model as different distribution channels chosen by a monopolistic manufacturer.

The diagnosis cost

The model assumes that the expert can identify the correct service without incurring any cost. The justification for assuming zero diagnosis cost is that, if the expert finds it profitable to refrain from providing diagnosis to some consumers when diagnosis costs are zero, then, *a forteriori,* she will do so with positive diagnosis costs. So, in studying price discrimination, there is no loss of generality in this assumption. Assuming zero diagnosis costs has also expositional advances: With a strictly positive diagnosis cost we would have to define – and concentrate on – parameter constellations for which performing a serious diagnosis and providing the diagnosed treatment is the efficient policy.[33] This would complicate the analysis without altering the results.

Consumers' sunk cost

The basic model of Section 2 assumes that each consumer incurs a sunk cost $c \geq c_2 - c_1$ if he visits the expert independently of whether he is actually treated or not.[34] This assumption has the effect that a consumer's option to reject a recommended treatment doesn't impose a binding constraint on the expert's maximization problem. To see this, consider the basic model of Section 2 and suppose that the *ex post* participation constraint is not binding. Then, the maximal profit the expert can realize from serving a type t consumer with equal mark-up prices is $v - (1-t)c_1 - tc_2$ (otherwise the consumer would refrain from visiting the expert) while equal mark-up prices require $p_2 - c_2 = p_1 - c_1$. Thus, $p_2 = c_2 + v - (1-t)c_1 - tc_2 = v + (1-t)(c_2 - c_1)$. Now, from $c \geq c_2 - c_1$ it follows that $p_2 < \hat{v} = v + c$ so that the *ex post* participation constraint is indeed not binding. If we allowed for $c < c_2 - c_1$ then there could be parameter constellations for which the restriction $p_2 \leq \hat{v} = v + c$ becomes binding. In this case, the largest surplus that could be extracted with an equal mark-up tariff would be $(1-t)[v+c-(c_2-c_1)]+t(v+c)$ leading to a profit of $(1-t)[v+c-(c_2-c_1)-c_1]+t(v+c-c_2) = v+c-c_2$. This profit is larger than the profit attainable with an overtreatment tariff (which is $v - c_2$) whenever $c > 0$. Thus, allowing for $c < c_2 - c_1$ would complicate the analysis without providing different results (provided $c > 0$).

An alternative interpretation of the model

The model and the analysis were developed in the context of a monopolistic expert who second degree price-discriminates among consumers via a menu of tariffs. Here, we argue that it is possible to re-interpret this scenario and the results we obtained in terms of a monopolistic manufacturer who uses different distribution channels as an instrument of price discrimination. A simple re-interpretation of the results of the basic model – where the expert posts an equal mark-up tariff to skim-off low-cost consumers and an

[33] For instance, in the basic version of the model – where there are only two types of problem and two types of treatment – performing a serious diagnosis and providing the diagnosed treatment is the efficient policy if and only if the diagnosis cost d satisfies $d \leq \min\{(1-g_2)(c_2 - c_1), g_2(\hat{v} - c_2 + c_1)\}$.

[34] In Subsection 4.2 – where consumers differ in their gross valuation $\hat{v}$ but have the same g_2 – the relevant condition is $c \geq (1-g_2)(c_2 - c_1)$.

overtreatment tariff to serve the rest – is that the monopolistic manufacturer sets up two types of stores, specialist outlets through which she distributes the entire assortment and where the qualifications and incentives of the sales personnel are such that they diagnose customers' needs and suggest the right equipment; and discount outlets through which she distributes only the major treatment. This story is not very realistic, however. A more elaborate model would have a monopolistic manufacturer that distributes her products through a competitive retail stage. The simplest version of such a model would have two types of retailers, expert shops with highly qualified sales personnel, and discounters. Suppose that there are at least two retailers of each of these two types in the market. Further suppose that the manufacturer's products are only a small part of a much larger number of products handled by a typical retailer so that the manufacturer can treat the characteristics of the retail stage as given.[35] Then the manufacturer can mimic the single-expert behavior by the choice of different distribution channels. To see this, suppose that the manufacturer wants to skim-off low-cost consumers via the Δ_{12}-tariff ($p_1 = c_1 + \Delta_{12}$, $p_2 = c_2 + \Delta_{12}$) and serve the rest of the market with the Δ_{02}-tariff ($p_1 < c_1 + \Delta_{02}$, $p_2 = c_2 + \Delta_{02}$). How can she do this? She offers the two treatments at wholesale prices $w_1^e = c_1 + \Delta_{12}$ and $w_2^e = c_2 + \Delta_{12}$ to at least two expert shops and she offers only the major treatment at the wholesale price $w_2^d = c_2 + \Delta_{02}$ to at least two discounters. With at least two discounters carrying the major treatment, market equilibrium yields $p_2^d = w_2^d$ by the usual price-undercutting argument. And expert retailers? Were an expert to post prices violating the equal mark-up rule, consumers would become suspicious; they would correctly infer that the expert will either always recommend the major treatment (if $p_1^e - w_1^e < p_2^e - w_2^e$), or always recommend the minor one (if $p_1^e - w_1^e > p_2^e - w_2^e$), and they would adjust their willingness to pay accordingly. So, experts cannot gain from cheating. Consequently, at least two experts post tariffs that induce non-fraudulent behavior. With tariffs that induce non-fraudulent behavior, Bertrand competition again yields prices such that underbidding yields losses and charging more implies a loss of customers; that is, $p_1^e - w_1^e = p_2^e - w_2^e = 0$.

Note that here again the sunk cost c plays an important role in making the story consistent: For small enough c ($c < p_2^e - p_2^d$) experts would become vulnerable to competition by discounters. Why? Because discounters would then be able to attract consumers who have learned from an expert that

[35] Marvel and McCafferty (1984) make a similar assumption in a different context.

they need a high quality treatment. To avoid this, the manufacturer would have to reduce w_2^e accordingly. The rest of the story is the same as that in the basic model (discussed in the previous paragraph): The profitability of price discrimination would be reduced but price discrimination would remain profitable as long as $c > 0$.[36]

6 Concluding Remarks

Research on credence goods markets typically assumes that consumers are homogeneous. The present article has studied the consequences of dropping this assumption in a model where an expert / a manufacturer has some degree of market power in serving the market. With heterogeneous consumers and market power price discrimination may emerge in equilibrium. We have shown, that in the case of experts markets, price discrimination proceeds along the dimension of quality of advice offered. High quality diagnosis and appropriate treatment is sold to the most profitable market segment only. Less profitable consumers are induced to demand a procedure without a serious diagnosis. If consumers differ in the probability of needing different interventions then low-cost consumers get high quality diagnosis and appropriate treatment while high-cost consumers are potentially *overtreated*; that is, they are induced to demand high quality equipment without a serious diagnosis. By contrast, under heterogeneity in the valuation of a successful intervention high-valuation consumers get efficient diagnosis and appropriate treatment, while low-valuation ones are potentially *undertreated*; that is, they are induced to demand low quality equipment independently of the severity of their problem.

While the equilibrium behavior outlined in the present paper is obviously an abstraction and it is probably impossible to point out an industry that features exactly this kind oft price discrimination, our results identify an element that may be present in the conduct of many credence goods industries. We discussed the evolution of the IT industry in the introduction. Other credence goods markets with similar characteristics – large heterogeneity in consumers' valuation for a successful match in the beginning; large heterogeneity in the product quality with product maturation – featured a similar evolution. Examples are fitness equipment, for example treadmills. At the beginning advanced machines where only available through specialized stores that also

[36] Here, remember that in the optimal menu of Section 3 $\Delta_{02} = v$.

offered advice in choosing the right equipment, whereas simpler equipment was sold at warehouse outlets. Nowadays, advanced equipment is available from discounters. Similarly, better makes of digital cameras were first sold by expert stores only, whereas now, also discount outlets offer advanced digital photo equipment. Digital hearing aids, stereo equipment, functional sports ware and even eyeglasses experienced a similar history.

References

[1] Acharyya, R. (1998), "Monopoly and Product Quality: Separating or Pooling Menu?", *Economics Letters* **61**, 187-194

[2] Alger, I. and Salanié, F. (2002), "A Theory of Fraud in Experts Markets", Mimeo, Department of Economics, Université de Toulouse

[3] Darby, M. and Karni, E. (1973), "Free Competition and the Optimal Amount of Fraud", *Journal of Law and Economics* **16**, 67-88.

[4] Dulleck, U. and Kerschbamer, R. (2003), "On Doctors, Mechanics and Computer Specialists. Or, Where are the Problems with Credence Goods", Working Paper, Department of Economics, University of Vienna

[5] Emons, W. (1997),"Credence Goods and Fraudulent Experts", *Rand Journal of Economics **28**,* 107-119.

[6] Emons, W. (2001), "Credence Goods Monopolists", *International Journal of Industrial Organization* **19**, 375-389.

[7] Gabszewicz, J. J. and Wauthy, X. Y. (2002), "Quality Underprovision by a Monopolist When Quality is Not Costly", *Economics Letters* **77**, 65-72

[8] Glazer, J. and McGuire, T. (1996), "Price Contracts and Referrals in Markets for Services", Working Paper No 10/96, Institute of Business Research, Tel Aviv University.

[9] Marvel, H. and McCafferty, S. (1984), "Resale Price Maintenance and Quality Certification", *RAND Journal of Economics* **15**, 346-359.

[10] Mussa, M. and Rosen, S. (1978), "Monopoly and Product Quality", *Journal of Economic Theory* **18,** 301-317

[11] Nelson, P. (1970), "Information and Consumer Behavior", *Journal of Political Economy* **78,** 311-329.

[12] Pesendorfer, W. and Wolinsky, A. (2003), "Second Options and Price Competition: Inefficiency in the Market for Expert Advice", *Review of Economic Studies* **70**, 417-437.

[13] Pitchik, C. and Schotter, A. (1987), "Honesty in a Model of Strategic Information Transmission", *American Economic Review* **77**, 1032-1036; Errata (1988), *AER* **78**, 1164.

[14] Pitchik, C. and Schotter, A. (1993), "Information Transmission in Regulated Markets", *Canadian Journal of Economics* **26**, 815-829.

[15] Richardson, H. (1999), "The Credence Good Problem and the Organization of Health Care Markets", mimeo, Department of Economics, Texas A&M University.

[16] Taylor, C. (1995), "The Economics of Breakdowns, Checkups, and Cures", *Journal of Political Economy* **103**, 53-74.

[17] Wolinsky, A. (1993), "Competition in a Market for Informed Experts' Services", *Rand Journal of Economics* **24**, 380-398.

[18] Wolinsky, A. (1995), "Competition in Markets for Credence Goods", *Journal of Institutional and Theoretical Economics* **151**, 117-131.