

Dulleck, Uwe; Johnston, David W.; Kerschbamer, Rudolf; Sutter, Matthias

Working Paper

The good, the bad and the naive: Do fair prices signal good types or do they induce good behaviour?

IZA Discussion Papers, No. 6491

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Dulleck, Uwe; Johnston, David W.; Kerschbamer, Rudolf; Sutter, Matthias (2012) : The good, the bad and the naive: Do fair prices signal good types or do they induce good behaviour?, IZA Discussion Papers, No. 6491, Institute for the Study of Labor (IZA), Bonn, <https://nbn-resolving.de/urn:nbn:de:101:1-2012080812340>

This Version is available at:

<https://hdl.handle.net/10419/58978>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

IZA DP No. 6491

**The Good, the Bad and the Naive:
Do Fair Prices Signal Good Types or
Do They Induce Good Behaviour?**

Uwe Dulleck
David Johnston
Rudolf Kerschbamer
Matthias Sutter

April 2012

The Good, the Bad and the Naive: Do Fair Prices Signal Good Types or Do They Induce Good Behaviour?

Uwe Dulleck

Queensland University of Technology

David Johnston

Monash University

Rudolf Kerschbamer

University of Innsbruck

Matthias Sutter

*University of Innsbruck,
University of Gothenburg and IZA*

Discussion Paper No. 6491

April 2012

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0

Fax: +49-228-3894-180

E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

The Good, the Bad and the Naive: Do Fair Prices Signal Good Types or Do They Induce Good Behaviour?*

Evidence on behavior of experts in credence goods markets raises an important causality issue: Do “fair prices” induce “good behavior”, or do “good experts” post “fair prices”? To answer this question we propose and test a model with three seller types: “the good” choose fair prices and behave consumer-friendly; “the bad” mimic the good types’ price-setting, but cheat on quality; and “the naive” fall victim to a projection bias that all sellers behave like the bad types. OLS, sample selection and fixed effects regressions support the model’s predictions and show that causality goes from good experts to fair prices.

JEL Classification: C91, L15, D82, D40

Keywords: credence goods, experts, pricing, experiment, other regarding preferences, signalling, projection bias

Corresponding author:

Rudolf Kerschbamer
Department of Economics
University of Innsbruck
Universitätsstraße 15
6020 Innsbruck
Austria
E-mail: rudolf.kerschbamer@uibk.ac.at

* We gratefully acknowledge financial support from the Australian Research Council (ARC) through grant number DP110103653, the Austrian Science Fund (FWF) through grant number P20796 and from the Austrian National Bank (OeNB Jubiläumsfonds) through grant number 13602. Part of this research was conducted while the third author was visiting the Queensland University of Technology and he would like to thank its School of Economics and Finance for its great hospitality.

1 Introduction

Trades involving credence goods are affected by asymmetric information, because expert sellers are better informed than their customers about the quality that yields the highest surplus from trade (Darby and Karni, 1973; Dulleck and Kerschbamer, 2006). As a consequence, many markets for credence goods are plagued by inefficiencies. Common examples for credence goods are health services, where the doctor knows better than the patient which disease the latter has and which treatment is appropriate (e.g., Gruber and Owings, 1996; Gruber et al., 1999; Iizuka, 2007); car repair services, where the mechanic knows more about the type of service the vehicle needs than the owner (e.g., Hubbard, 1998; Wolinsky, 1993, 1995); and taxicab rides in an unknown city, where the driver is better informed about the shortest route to the destination than the passenger (e.g., Balafoutas et al., 2011). All studies mentioned here find that inefficiencies result at least partly from material incentives for sellers to cheat on customers. However, the outlook is not too bleak as the same studies also reveal that a significant number of experts provide efficient service even when material incentives are misaligned.

Recent experimental evidence also suggests that a non-negligible fraction of experts ‘do the right thing’ even when material incentives are strongly distorted. For instance, for an environment without institutional safeguards against (and therefore strong material incentives for) fraud, Dulleck et al. (2011) find that almost 30% of experimental experts provide the appropriate quality throughout their experiment. They also present evidence showing that “fair” prices – that is, prices that split the gains of trade equally between the two parties of a transaction – are much more frequently observed than predicted, and that they are atypically often associated with efficient provision behavior despite the fact that they provide material incentives for (inefficient) undertreatment. It is exactly this observation that raises an interesting causality issue: Do “fair prices” induce “good behavior”, or do “good experts” post “fair prices”? The answer to this question has important practical implications, for instance for regulation policy: If fair prices induce good behavior then price regulation might be a cheap and effective means to increase efficiency in markets for credence goods; however, if good experts post fair prices then price regulation may actually reduce efficiency, as the prices posted by expert sellers in unregulated markets contain valuable information for customers about the type of seller they face, information that is destroyed by regulation.

This paper addresses this causality issue, that is, the question whether “fair prices” induce “good behavior”, or “good experts” post “fair prices”? It does so by first presenting a model with heterogeneous sellers, and then running new experiments to test the main implications of the model. Motivated by existing empirical and experimental evidence our model features three types of experts, “the good”, “the bad” and “the naive”.

“The good” receive an extra utility from solving the consumer’s problem, for instance, because they value efficiency, or because they care for the well-being of consumers (Charness and Rabin, 2002). Consistent with the evidence reported above we assume that this extra utility is large enough to prevent the good experts from providing insufficient quality, even if such behavior is against their monetary self-interest.

“The bad” are rational own-money maximizing experts who take advantage of customers whenever there are material incentives for doing so. Since our bad experts are just the textbook agents, they (probably) need no further explanation or motivation.

“*The naive*” have the same (own-money-maximizing) preferences as the bad, but differ in their expectations about the preferences of others. Specifically, our model assumes that the naive have a pronounced tendency to overestimate the degree of alignment of their own preferences with those others have.¹ Experimental evidence in support of this assumption comes from a recent study by Iriberry and Rey-Biel (2010). They use the observed behavior of subjects in modified dictator games to classify them into different archetypes of distributional preferences and then investigate whether the existing heterogeneity in distributional preferences is correctly anticipated by subjects. They find that selfish individuals cannot conceive others being non-selfish, while the beliefs of cake-size maximizers are closest to the actual distribution of others’ preferences. An explanation for the fact that selfish subjects have a more pronounced tendency than others to assume that their preferences are kind of ‘normal’ and that others also think the same way as they do comes from the psychological literature on the “projection bias”. According to the theories of Sigmund Freud a person suffering from a projection bias projects his own undesirable thoughts, motivations, desires, feelings, and so forth onto someone else in order to avoid the discomfort of consciously admitting a personal fault.² In our context, the projection bias serves the function of relieving selfish experts from the psychological strain they would feel when breaking a code of conduct others respect, since it seems plausible that providing the appropriate quality is the socially approved behavior in credence goods markets.

In our model the projection bias of some selfish experts leads to the existence of a signalling equilibrium where experts posting fair prices are more likely to provide efficient service than those posting different prices. Specifically, in the proposed equilibrium “the good” choose fair prices and behave in a consumer-friendly manner; “the bad” mimic the good in their price-setting behavior but take advantage of customers in their provision and charging decisions; and “the naive” randomize in their price-setting behavior and take advantage of consumers in their provision and charging decisions if they find a victim. Customers understand the signalling value of prices and are hence more prone to accept if fair prices are posted.

Our model implies a causal effect from other-regarding preferences towards pricing behavior and not vice versa. Pure theories of other-regarding behavior – for example, the prominent models of inequity aversion by Fehr and Schmidt (1999) and Bolton and Ockenfels (2000) – are consistent with both directions of causality, i.e. setting prices that are particularly fair incentivizes inequity averse experts to provide efficient service, as well as inequity averse experts posting such prices more frequently. We confirm the causal unidirectionality empirically, by comparing experimental conditions in which prices are endogenous, i.e. chosen by the expert sellers, to conditions where prices are exogenous, i.e. randomly set

¹The tendency of persons to overestimate the degree of agreement that others have with them is often termed “false consensus effect” or “consensus bias”. This cognitive bias was first examined by psychologists – Ross et al. (1977) and Mullen et al. (1985) are classical references – but it has also attracted attention in economics, see Selten and Ockenfels (1993) or Charness and Grosskopf (2001).

²The psychological literature provides evidence for the existence of (attributive) projection bias (Holmes 1978, Sherwood 1981) which is characterized as a “defence mechanism to protect our self-esteem. If we break a work rule, projection bias justifies this infraction by claiming that ‘everyone does it’.” (McShane and Travaglione 2003, p. 87). Economist often use the term in a slightly different meaning, see Loewenstein et al. (2003), for instance.

by the experimental software. Estimates from OLS and sample selection regression models suggest that with endogenous prices the fact whether a transaction takes place under a "fair" or an "unfair" price-vector significantly affects both the probability with which the consumer accepts and the probability that she receives efficient service; however, those effects disappear, as predicted, when controlling for expert type with fixed effects, or by exogenously imposing prices. Thus, it is not the case that "fair prices" induce "good behavior", instead "good experts" frequently choose "fair prices".

With respect to the experimental literature, our paper is most closely related to Dulleck et al. (2011). Indeed, the current paper takes two of their 16 experimental treatments with endogenous prices as the starting point and adds two new ones where prices are exogenously fixed. While Dulleck et al. (2011) discuss the facts that experimental experts are heterogeneous in their price-posting, provision and charging behavior and that experimental experts frequently post fair prices, they do not address the endogeneity problem ("Do 'fair prices' induce 'good behavior', or do 'good experts' post 'fair prices'?"). By contrast, for the current paper, the observation that fair prices are associated with consumer-friendly behavior is central and the endogeneity issue is our main research question.³ Somewhat less directly related are three recent papers by Huck et al. (2006, 2007, and 2010). They use a binary version of the well-known trust game (by Berg et al. 1995) as a stylized representation of a market for experience goods and show that those goods are more efficiently provided when sellers can build up reputation than when they can not, and when there is competition in the market than when there is not. Related to our paper is their finding that independently of the market conditions (monopoly vs. oligopoly; fixed prices vs. flexible prices; reputation building possible vs. reputation building impossible) a significant share of sellers behave in a consumer-friendly way, despite the presence of strong material incentives for providing the low quality throughout. Huck et al. (2006, 2007, and 2010) do not study causality issues from price-setting behavior to sellers' types, however.

Regarding theory the present paper is related to two strands of literature: First, and most obviously, the present paper is related to the theoretical literature on efficiency and honesty in markets for credence goods (see Dulleck and Kerschbamer, 2006, for a survey). Secondly, the paper is related to a recent study by Andreoni and Bernheim (2009) on 'audience effects'. In their dictator game model, dictators have private information about their own preferences and like to be perceived by others as "fair". Specifically, dictators care about their own material payoff, about the fairness of the final allocation (where fairness is judged by the extent to which the outcome departs from the egalitarian distribution) and about their 'social image' (modelled as the audience's perception of the dictator's fairness). The asymmetric information on preferences together with the fact that the better informed party moves first gives rise to a signaling game wherein the dictator's choice affects the audience's inference about the dictator's taste for fairness. In the focal equilibrium of this game many dictators

³The present paper also shares a treatment with an unpublished working paper by Kerschbamer et al. (2009). While the main contribution of that paper is the identification of experts' distributional preference types from their provision behavior under exogenously fixed prices, the main focus of the current paper is the experts' price-posting behavior under endogenous prices and the question whether prices contain information about the price-posting seller's type. To answer this question, we take the experiments with exogenous prices as a control; by contrast, for Kerschbamer et al. (2009) the experiments with exogenous prices are the main instrument and the endogenous-price treatments are only discussed as a motivation.

voluntarily cede exactly half of the pie and no one gives either more or slightly less than half of the pie. Our model shares the feature that the 50-50 split is used as a signalling device and our experimental findings can be interpreted as providing support for the authors' conjecture that their theory "... that accounts for the 50-50 norm in the dictator game [...] will prove applicable more generally" (Andreoni and Bernheim 2009, p. 1607). An obvious difference to their model is our assumption that a fraction of the selfish experts falls victim to a projection bias (they are unaware of the signalling content of the 50-50 norm). Another difference is that in our model the recipient's inference matters to the sender only because it affects the probability of acceptance.

The rest of the paper is organized as follows. The next section introduces our simple theoretical model of credence goods featuring heterogeneous experts, some of them suffering from a projection bias. Section 3 presents the experimental design as well as three main hypotheses following from the theoretical model. Section 4 presents our main empirical results. Section 5 concludes.

2 The Model

We take the simplest version of Dulleck and Kerschbamer's (2006) model of credence goods as our starting point. In this model there are two players, an expert (he) and a consumer (she). The consumer has a problem t that is with ex ante probability h major ($t = t^h$) and with probability $1 - h$ minor ($t = t^l$). The consumer knows that she has a problem but does not know how severe it is; she only knows the ex ante probabilities of the two severities. If the consumer decides to visit the expert, the latter finds out the severity of the problem by performing a diagnosis. He then provides a service c that is either of high quality ($c = c^h$) or of low quality ($c = c^l$). The two qualities come at different cost for the expert, the cost of the high quality is c^h , the cost of the low quality is c^l , with $c^h > c^l$.⁴ The high quality service solves both severities of the problem, while the low quality is only sufficient for the minor problem. If the consumer's problem is solved, she receives a value of $v > 0$, otherwise she receives a value of zero.

The sequence of moves in our game is as follows: At the start the expert posts two prices, a price p^h for the high-quality service and a price p^l for the low-quality service, with $p^h \geq p^l$. After observing the two prices the consumer decides whether to visit the expert or not. If the consumer decides against a visit, the expert and the consumer both receive a material payoff of o as their outside option. If the consumer decides to visit the expert, the expert observes the severity of the consumer's problem. He then provides one of the two qualities (both qualities are feasible for both severities of the problem – in the language of Dulleck and Kerschbamer, 2006, there is *no liability*), and charges one of the two posted prices (both prices are feasible in the baseline game independently of the quality provided – in the language of Dulleck and Kerschbamer, 2006, there is *no verifiability* in the baseline game). After the provision and charging decision of the expert the game is over, and material payoffs are realized. The material payoff of the consumer for the case where she has visited the expert and her problem has been solved is v minus the price charged by the expert; if the consumer has visited the expert and received an insufficient service, her material payoff is zero minus

⁴For simplicity both the quality provided and the cost of the quality provided are denoted by c .

the price paid. The material payoff of the expert in case of interaction is the difference between the price charged (either p^h or p^l) and the cost of the treatment provided (either c^h or c^l). Throughout we assume that the value of v is such that $v - hc^h - (1 - h)c^l > 2o$, but $(1 - h)v - c^l < 2o$. While the former inequality means that solving the consumer's problem is efficient from an ex ante perspective, the latter implies that in the standard benchmark – where both players are rational, risk-neutral and only interested in their own material payoff and where this fact is common knowledge – the market breaks down, just as Akerlof's (1970) market for lemons.⁵

Motivated by the empirical and experimental evidence discussed in the introduction we depart from the default assumption in the credence goods literature that all experts are rational, risk-neutral and exclusively interested in their own material payoff by assuming that an expert is either "good" (type g), or "bad" (type b), or "naive" (type n). Good experts receive an extra utility from solving the consumer's problem and we assume that this extra utility is large enough to prevent them from providing c^l when c^h is required to solve the consumer's problem. For the rest, good experts are just like "conventional experts", they are rational and risk-neutral, and they care for their own material payoff.⁶ Bad experts are just "standard" agents. They are rational, risk-neutral and care only for their own material payoff, and they have correct beliefs regarding the distribution of expert types in the population. Naive experts have the same preferences as bad experts, the only difference to them is that they cannot perceive that there might be g -type experts in the market. The relative frequencies of the three types of sellers are denoted by α_g , α_b and α_n , respectively, with $\alpha_n = 1 - \alpha_g - \alpha_b$. We assume that those frequencies are known by consumers, as well as by g - and b -type experts, but that n -type experts believe that $\alpha_n = 1$.

If $\alpha_g = 0$ –that is, if all sellers are exclusively interested in their own material payoffs– then only low quality is provided and the high price is charged, just as in the benchmark. We assume that even with $\alpha_g > 0$ this is the equilibrium expected by n -type sellers. That is, whenever an n -type seller posts any price-vector with $p^h > (1 - h)v - o$ he expects that no customer will accept his offer and whenever he posts a price-vector that violates this condition he expects to earn less than the value of his outside option. Thus, posting any price-vector with $p^h > (1 - h)v - o$ is fine for him – it is in this sense that no price-vector has a particular meaning to an n -type seller. By contrast, for g -type sellers posting a price-vector with $p^h > (1 - h)v - o$ is meaningful if they expect customers to accept because they know about g -type sellers' existence.

In the following we characterize an equilibrium where the price-vector with $p^h = (v + c^h)/2$ and $p^l = (v + c^l)/2$ serves as a signalling device. This price-vector is the most frequently posted price-vector in the baseline treatments of the experiments reported by Dulleck et al. (2011) and it has the characteristic that under the assumption that sellers provide the appropriate quality and charge the price of the quality provided, the surplus is split half way independently of whether the consumer has the minor or the major problem. This makes

⁵To see that the assumption $(1 - h)v - c^l < 2o$ implies market break down in the standard benchmark note that a rational own-money maximizing expert will always provide c^l and always charge p^h under each price vector. Anticipating this behavior, the consumer is only prepared to visit the expert if $p^h \leq (1 - h)v - o$. But such a low p^h is unattractive for the expert because it implies $p^h - c^l < o$.

⁶Formally, good experts are assumed to maximize their material payoff under the constraint that the consumer's problem is solved.

this price-vector focal. It is also focal for other reasons, for instance, because it would be the outcome of Nash bargaining under full information and bilateral monopoly.⁷ Assuming that players are able to coordinate their expectations (possibly because the focality of the equal split helps them to coordinate) we claim:⁸

Proposition 1 *If the share of g -type sellers is large relative to the share of b -type sellers then there exists a Perfect Bayesian Equilibrium where*

- *all g - and b -type sellers post prices $p^h = \frac{v+c^h}{2}$ and $p^l = \frac{v+c^l}{2}$, while n -type sellers randomize over all possible price-vectors with $p^h > (1-h)v - o$ using a continuous distribution function without a mass-point;*
- *g -type sellers always provide the appropriate quality and charge p^h , b - and n -type sellers always provide low quality and charge p^h ;*
- *customers choose to interact if the seller posts a price-vector that has either $p^h = \frac{v+c^h}{2}$ and $p^l = \frac{v+c^l}{2}$, or $p^h \leq (1-h)v - o$ and they refrain from interaction if any other price-vector is posted; and*
- *customers believe that the expert is of type g with probability $\frac{\alpha_g}{1-\alpha_n}$ and of type b with probability $\frac{\alpha_b}{1-\alpha_n}$ if they observe a price-vector with $p^h = \frac{v+c^h}{2}$ and $p^l = \frac{v+c^l}{2}$ and they believe that the expert is of type n for sure if they observe any other price-vector.*

Proof. As is easily verified, all types of the expert behave optimally given their preferences and their beliefs about the behavior of consumers. Furthermore, customers' beliefs are consistent with their initial beliefs and experts' equilibrium behavior. Given their beliefs, acceptance is optimal for consumers iff $\alpha_g(v - p^h) + \alpha_b[(1-h)v - p^h] \geq (1 - \alpha_n)o$. Simplifying this inequality and replacing p^h by $\frac{v+c^h}{2}$ gives $\frac{\alpha_b}{1-\alpha_n} \leq \frac{(v+c^h)/2-o}{hv}$, which confirms that the share of g -type sellers has to be sufficiently large compared to the share of b -type sellers. ■

It is important to note that naive experts are needed for prices to contain information – without them the equilibrium would be pooling. As mentioned in the introduction there exists convincing experimental evidence indicating that selfish subjects have a more pronounced tendency to fall victim of a (false) consensus effect; that is, selfish subjects are, on average, more likely than others to assume that their preferences are ‘normal’ and that others also think the same way they do.

⁷Evidence supporting the claim that a 50-50 division is focal even in settings where one party unilaterally determines the allocation of a pie comes from dictator-game studies. For instance, in the continuous dictator games investigated by Andreoni and Miller (2002), a significant fraction of subjects share equally despite the material cost to themselves. Furthermore, as discussed in the introduction, Andreoni and Bernheim (2009) provide a theoretical model based on inequality aversion supporting such a split. They cite several studies showing that fifty-fifty sharing is common in the context of joint ventures among business firms, partnerships among professionals, share tenancy in agriculture, and bequests to children. They also provide evidence indicating that equality is a frequent outcome of negotiation and conventional arbitration in the field.

⁸In our model the fifty-fifty price-vector is, of course, not the only price-vector that has the potential to act as a signalling device. Later – in Section 4 – we will empirically ask whether other price-vectors are used in our experiments as a signalling device as well.




Figure 1: The Credence Goods Game Form

3 Experimental Design and Hypotheses

In this section we first introduce the experimental design (which closely resembles the basic credence goods model introduced in the previous section) and the parametrization. Then we present three hypotheses that follow from the theoretical model presented above. Each hypothesis describes the predicted behavior for one of the three stages of the game, that is, the price-posting stage, the interaction-decision stage, and the provision-and-charging stage.

We employ a 2x2 factorial design, varying the following two factors: Regarding institutional background we have treatments with verifiability (*WV*) and treatments without verifiability (*NV*), regarding prices we have treatments with endogenous price-setting (*endo*) and treatments with exogenous prices (*exo*). This gives the four treatments *WV-endo*, *WV-exo*, *NV-endo*, *NV-exo*. *NV-endo* corresponds exactly to the game described in the previous section and displayed in Figure 1. In this treatment experts are free to charge either p^h or p^l , independently of whether they have provided c^h or c^l . *WV-endo* differs from *NV-endo* only in that it forces experts to charge the price of the quality provided – that is, in this treatment the last stage of the game displayed in Figure 1 is degenerate.⁹ The two *exo* treatments differ from their *endo* counterparts only by the fact that prices are fixed exogenously – that is, in those treatments the first stage of the game displayed in Figure 1 is degenerate. Table 1 summarizes the different treatments and includes already the number of subjects and the relative frequency of interaction in the experiment, to which we will refer later.

⁹A real world interpretation is that consumers are able to observe and verify *ex post* the quality they received (without knowing, however, whether this quality is the appropriate one) – see Dulleck and Kerschbamer (2006) for a discussion. While standard theory predicts that verifiability is important for efficiency in markets for credence goods, the experimental results reported by Dulleck et al. (2011) show that it has at best a minor effect.

Table 1: Treatments, Number of Subjects and Percentage of Interaction by Customers

	NO Verifiability (<i>NV</i>)	With Verifiability (<i>WV</i>)
Prices Endogenous (<i>endo</i>)	<i>NV-endo</i>	<i>WV-endo</i>
# subjects	96	88
% interact	44.9	50.0
Prices Exogenous (<i>exo</i>)	<i>NV-exo</i>	<i>WV-exo</i>
# subjects	64	128
% interact	50.4	63.0

Note: Each subject in each treatment plays 16 rounds; % interaction refers to the average rate of acceptance by consumers, where the average is taken over subjects and over rounds.

In all treatments the customer’s probability of needing the high quality is set to $h = 0.5$, and the value of sufficient quality is set to $v = 10$. The costs of providing the low, respectively high, quality is $c^l = 2$, respectively $c^h = 6$. In the *endo* treatments, experts are asked to post prices p^l and p^h from the finite set $\{1, 2, \dots, 10, 11\}$, respecting the restriction $p^l \leq p^h$. In the *exo* treatments, the price-vector (p^l, p^h) was chosen randomly and with equal probability from the set $\{(3,8), (4,8), (5,8), (6,8), (7,8)\}$. Four of these price-vectors are among the five most popular price-vectors in the *endo* treatments (see Table 2). The value of the outside option for the case where no trade takes place is set to $o = 1.6$.

In all treatments we implement matching groups of eight subjects, with four subjects in the role of customers and four subjects in the role of sellers. Role assignment is random at the beginning of the experiment and fixed for all 16 periods of the experiment. In order to avoid reputation building customers are randomly re-matched to sellers within each matching group after each period.

All experimental sessions were computerized using zTree (Fischbacher 2007). Recruiting was done via ORSEE (Greiner 2004). Sessions started with an extensive description of the game and only after subjects had correctly answered a set of control questions. All parameters as well as the matching procedure were made common knowledge to all participants by reading them aloud (see the experimental instructions in the appendix). The average session length was 1.5 hours, and subjects earned on average 15 Euro.

The experimental design as well as the parameters are in line with the assumptions made in the theoretical part. The price-vector $(p^l, p^h) = (6, 8)$ is the one discussed in the theory section and our “theory” would predict that this price-vector is disproportionately often associated with efficient service and that it is therefore accepted by customers. Naturally, we do not expect our model to predict experimental behavior perfectly and the experimental results indeed indicate that a wide range of price-vectors are chosen by the experimental experts in the treatments with endogenous prices (see Table 2). Thus, we restrict ourselves to testing the following hypotheses which reflect, in our eyes, the qualitative messages of our model. The order of the hypotheses follows the backward induction logic inherent in the concept of Perfect Bayesian Equilibrium. That is, we start with experts’ provision and charging behavior, given that the posted price-vector has been accepted; next we consider customers’ acceptance decision; and only then experts’ price-posting choices.

Hypothesis 1 (Prices signal seller type) *With endogenous prices, experts posting the*

Table 2: Absolute Frequency of Experimental Price-Vectors in *NV-endo* and *WV-endo* (data from Dulleck et al., 2011)

		p^h										
		1	2	3	4	5	6	7	8	9	10	11
p^l	1					3	1	1	16	4		
	2		1	2	1	3	5	2	11	1	1	1
	3					15	21	31	21	6	1	
	4				11	6	21	40	101	9	1	
	5					12	33	78	90	28	6	1
	6						34	108	441	14	14	
	7							38	120	53	3	6
	8								22	11	5	1
	9									2	2	5
	10										2	5
	11											1

price-vector (6,8) are disproportionately more likely to provide the appropriate quality (in particular when the high quality is needed) than experts posting other prices. If prices are set exogenously, or if a control for seller type is added, price-vector (6,8) will no longer be associated with efficient service.

If the theoretical model was perfectly predictive of behavior within the experiment, appropriate quality would only be observed for the signalling price-vector (6, 8). In the current form of the hypothesis, however, we allow g -type sellers to make mistakes in posting prices, and consequently price-vector (6, 8) is hypothesized to be strongly associated with efficient service and other price-vectors are predicted to be strongly associated with inefficiency. If prices are fixed exogenously, g -, b - and n -type experts are equally likely to be allocated the price-vector (6, 8) and consequently this vector loses its signalling value. Similarly, price-vector (6, 8) loses its signalling value if the seller's type is controlled for, as it is the type (g , b or n) that determines both prices and efficiency.

Hypothesis 2 (Customers respond to price signals) *With endogenous prices, consumers confronted with the price-vector (6,8) are disproportionately more likely to trade than consumers exposed to other price-vectors. If prices are set exogenously, price-vector (6,8) will no longer be associated with a higher interaction rate.*

According to our theoretical model, consumers should only accept to be served when the expert has posted the signalling price-vector (6, 8). We allow customers to make mistakes in this regard, but would expect to see a large significant effect of price-vector (6, 8) on the probability of interaction. Again, if prices are set exogenously, price-vector (6, 8) loses its signalling value, and the effect of price-vector (6, 8) on the probability of interaction should disappear.

Hypothesis 3 (Sellers are of three types) *With endogenous prices, there are two groups of sellers, those posting the price-vector (6,8) with high frequency and those randomizing over all feasible price-vectors. Furthermore, with respect to provision behavior, sellers*

who choose the price-vector (6,8) frequently are of two types, one frequently providing appropriate quality and the other almost always providing low quality.

As usual, behavior within the experiment will depart from the one outlined in our theoretical model; in particular, subjects will experiment in the first few rounds. Furthermore, some subjects may behave ‘erratically’ throughout the experiment. In our empirical analysis we concentrate on the last ten rounds to minimize learning effects, and in these rounds we will look for the existence of three types: one randomizing over prices, two posting price-vector (6,8) frequently, and among the latter two types, one consistently delivering appropriate quality and the other consistently providing low quality.

Before directly addressing hypotheses 1-3, in Table 2 we present information on the frequency with which experimental experts in the *endo* treatments choose different price-vectors. Table 2 shows that sellers post price-vector (6,8) more frequently than any other price-vector, in fact in 30% of cases. The next most frequent price-vectors, (7,8), (6,7) and (4,8), are chosen approximately only one-quarter as often as (6,8).

4 Empirical Results

4.1 Hypothesis 1 (H1)

Hypothesis 1 (H1) states that experts who post the price-vector (6,8) are more likely to provide appropriate quality than experts who post other price-vectors. In addition, it is hypothesized that price-vector (6,8) is associated with efficient service solely because sellers who are more likely to provide appropriate quality (*g*-type experts) choose (6,8) frequently. We initially test H1 by using all 16 rounds of data from treatments with endogenous prices (*NV-endo* and *WV-endo*), and estimating efficient service regression models of the form

$$eff_{it} = \alpha_0 + \alpha_1 s_{it}^* + \alpha_2 X_{it} + \theta_i + \varepsilon_{it}, \quad (1)$$

where eff_{it} equals one if seller i in period t provides the appropriate quality (i.e. efficient service) and zero otherwise, s_{it}^* equals one if the seller posts the signaling price-vector (6,8) and zero otherwise, X_{it} is a vector of control variables that vary across periods, θ_i is seller’s unobserved type, and ε_{it} is a random error term. The vector of control variables includes an indicator for whether the treatment has verifiability, the period number (1 to 16), and an indicator for whether the customer had a major problem (implying high quality was needed). Linear regression models are used to estimate equation (1), and also subsequent equations, in preference to probit or logit regression models. Linear models are preferred primarily because the interpretation of parameter estimates is more straightforward; this is especially important in specifications that include fixed-effects (see Angrist and Pischke, 2009). It is particularly noteworthy, though, that the results are robust to the linearity assumption (Appendix Table A1 includes estimated effects from binary choice models).

The OLS regression estimates in column 1 of Table 3 suggest that sellers who post price-vector (6,8) are 13.7 percentage points more likely to provide appropriate quality than sellers posting alternative vectors – a large and statistically significant effect (p -value = 0.001). The only other significant coefficient in column 1 belongs to the indicator for a major problem. As expected, sellers are significantly less likely to provide an efficient service if the customer

Table 3: Estimated Effect of Price-Vector (6,8) on Seller’s Provision Decision

	Endogenous Price (<i>endo</i>)				Exogenous Price (<i>exo</i>)	
	(1)	(2)	(3)	(4)	(5)	(6)
vector (6,8)	0.137*** (0.040)	0.097** (0.041)	0.102*** (0.036)	-0.031 (0.046)	-0.006 (0.032)	-0.034 (0.035)
verifiability	-0.053 (-0.042)	-0.069 (0.046)	-0.073** (0.035)	-	-0.046 (0.049)	-0.046 (0.049)
period	-0.003 (0.003)	0.000 (0.004)	-0.001 (0.005)	-0.004 (0.004)	-0.002 (0.003)	-0.001 (0.003)
major problem	-0.514*** (0.051)	-0.515*** (0.051)	-0.534*** (0.033)	-0.500*** (0.054)	-0.251*** (0.041)	-0.256*** (0.041)
price p^l	-	-0.001 (0.016)	0.006 (0.016)	-0.015 (0.019)	-	0.023* (0.012)
price p^h	-	0.055*** (0.021)	0.060*** (0.019)	0.029 (0.022)	-	-
sample size	697	697	697	697	903	903

Note: Figures in columns 1, 2, 5 and 6 are coefficients from OLS regression models. Figures in column 3 are coefficients from a maximum likelihood sample selection model. Estimate of rho from the selection model is 0.076 with standard error 0.122. Figures in column 4 are coefficients from a linear fixed-effect regression model. Clustered standard errors shown in parentheses. *, ** and *** denote significance at .10, .05 and .01 levels.

requires high quality, as in this instance providing low quality is the dominant behavior for b - and n -type sellers. Verifiability and period number are both quantitatively unimportant for efficiency.

In column 2, we test whether the (6,8) estimate in column 1 is suffering from omitted variable bias. It is possible that prices and efficiency are related in a manner not captured by the signaling price-vector variable alone. If so, the estimated (6,8) effect will be biased. For example, suppose that in truth (6,8) is not a signalling price-vector, and that instead efficient service is driven by the magnitude of p^l and p^h : the higher the posted prices, the more likely the expert will provide appropriate quality. In this instance, the coefficient on price-vector (6,8) in specification (1) will be biased upwards, as (6,8) is a relatively "expensive" price-vector (only 7% of offered price-vectors have a higher p^h). To control for possibilities such as this, we include in specification (2) covariates representing p^l and p^h .¹⁰ We find that efficiency is strongly related to p^h – an increase in p^h by one unit increases the probability of appropriate quality by around 6 percentage points. However, the effect of including these additional controls on the (6,8) coefficient is small – sellers who offer price-vector (6,8) are estimated to provide efficient service 9.7 percentage points more often than sellers who offer alternative price-vectors.¹¹

¹⁰In order to capture potential interaction effects, we also estimated regression models with $p^l \times p^h$ included as an additional regressor. We do not present these results, however, as the interaction term was insignificant in all models, and since the estimated effect of vector (6,8) remained unchanged. Similarly, we have estimated models allowing for quadratic price terms. Again, the main results were unchanged.

¹¹The equivalent probit marginal effect estimate equals 12.1 percentage points (p -value = 0.012). See

Another potential issue is that efficient service is only observed if customers choose to interact, which occurs in 47% of occasions in the treatments with endogenous prices. If interaction is observed for a sample of sellers who are unobservably different from other sellers, our estimates could suffer from sample selection bias. Fortunately, in our set-up this is unlikely. Sellers and consumers are randomly and anonymously matched in our game, and the only decision sellers make before consumers decide on whether to interact with them, is the price-posting. Hence, the only (observable or unobservable) differences between sellers who interact and sellers who do not, are just the posted prices. Given we control for prices in our regression models, our coefficient estimates should therefore be unbiased. In any event, we can test for sample selection bias by jointly estimating efficiency equation (1) and interaction equation

$$int_{it} = \mathbf{1}(\beta_0 + \beta_1 s_{it}^* + \beta_2 X_{it} + \beta_3 Z_{it} + \theta_i + u_{it} > 0), \quad (2)$$

where int_{it} equals one if customer i in period t chooses to interact and zero otherwise, $\mathbf{1}(\cdot)$ is an indicator function, u_{it} is a random error term, and Z_{it} is a vector that includes an indicator for whether the customer chose to interact in period $t - 1$ and the customer's profit in period $t - 1$. It is assumed that ε_{it} and u_{it} follow a bivariate normal distribution with correlation coefficient ρ . This classic sample selection model is identified by the variables Z_{it} that are included in equation (2), but which are excluded from equation (1). These excluded variables – customer interaction and profit in $t - 1$ – significantly affect the customer's interaction decision, but can not affect the seller's provision decision because the seller has no information about the customer. Estimates of the efficiency equation from the sample selection model are given in column 3 of Table 3.¹² As predicted, the estimates are largely unchanged. Furthermore, the correlation coefficient ρ is statistically insignificant (p -value = 0.533), implying that sellers who interact are not unobservably different from sellers who do not.

The results in columns (1), (2) and (3) show that customers are more likely to receive appropriate quality if they interact under the price-vector (6,8) than if they interact under other price-vectors. However, it does not show that there are no other 'special' price-vectors; that is, price-vectors that are associated with atypically high frequencies of efficient service. We test for this possibility by sequentially replacing the (6,8) dummy variable in specification (2) with dummy variables representing other price-vectors for which we observe interaction (e.g. (3,5), (3,6), (3,7), etc.).¹³ The results from this exercise (available upon request) show that (6,8) is the only price-vector that has a significantly positive effect on efficiency.

To test the second part of H1 – that price-vector (6,8) is only important because it signals a seller's type – we use two approaches. First, we control for sellers' unobserved type by

Column (1) in Appendix Table A1.

¹²The indicator for whether the customer chose to interact in period $t - 1$ and the buyer's profit in period $t - 1$ (i.e. the exclusion restrictions) are both statistically significant. The associated z -statistics equal 3.88 and 7.79, respectively, and the joint-test of insignificance has a Chi-squared statistic of 63.06. Furthermore, both variables have the anticipated effects: buyers who chose to interact last period and buyers who profited last period are more likely to interact this period.

¹³All in all we estimate 30 OLS and probit regression models, each with a different price-vector dummy variable and the control variables included in column (2) of Table 3. Importantly, some price vectors were not tested, either because there was no observed interaction or because there was no variation in provision behavior (a necessary condition for probit models).

estimating a fixed-effect regression model. So long as the seller type is time-invariant, this model will control for the effect of seller type on the probabilities of choosing price-vector (6,8) and providing efficient service. The fixed-effects estimates are presented in column (4) of Table 3 and clearly show that price-vector (6,8) no longer affects the likelihood of efficient service.¹⁴ Our second approach is to re-estimate specifications (1) and (2) with data from the *exo* treatments (*NV-exo* and *WF-exo*). The rationale behind this approach is that prices lose their signaling value in these games, so only the incentive value remains. Again, the estimates confer with H1 – once the signal is removed from price-vector (6,8), its effect on efficiency disappears (see columns 5 and 6).

Columns (5) and (6) in Table 3 also show that the estimated coefficient on 'major problem' is reduced; from -0.500 in column (4) to -0.256 in column (6). One reason for the smaller 'major problem' coefficient is that prices in the *exo* treatments are higher, on average, than those in the *endo* treatments. If model (2) is re-estimated with the sub-sample of *endo* transactions that take place under price-vectors $3 < p^l < 7$ and $p^h = 8$, the estimated effect of 'major problem' drops to -0.386. Another reason for the smaller 'major problem' coefficient in columns (5) and (6) is that the probability of efficient service is significantly lower in *exo* games; exogenous pricing is estimated to reduce efficiency by 7.4 percentage points (p -value = 0.026).

In summary, the results in Table 3 strongly support H1. Moreover, the results appear robust to potential issues of omitted variable and sample selection bias. The results are also robust to the modelling choice: estimates from OLS, probit and logit regression models are very similar.

4.2 Hypothesis 2 (H2)

Hypothesis 2 (H2) states that customers are more likely to accept if the posted price-vector is (6,8) than if it is another price-vector. It also asserts that this relationship is driven by customers associating price-vector (6,8) with *g*-type sellers. To test H2 we follow a similar approach to the one employed in our analysis of H1. First, we use 16 rounds of data from the *endo* treatments, and estimate regression models of the form

$$int_{it} = \gamma_0 + \gamma_1 s_{it}^* + \gamma_2 X_{it} + \theta_i + v_{it}, \quad (3)$$

where int_{it} equals one if customer i in period t chooses to interact and zero otherwise, and X_{it} includes an indicator for verifiability, period number, and linear p^l and p^h effects. In some specifications the customer fixed-effect θ_i is assumed to be uncorrelated with s_{it}^* (i.e. in simple OLS regression models), and in other specifications this assumption is not made (i.e. in customer fixed-effect regression models). The estimates are presented in Table 4.

Estimates using data from the *endo* treatments are presented in columns (1) to (3). The OLS estimates in (1) suggest that sellers who post price-vector (6,8) are 19.9 percentage points more likely to have their offer accepted than sellers posting alternative vectors (p -value < 0.001). Importantly, this result is not dependent upon our modelling approach. If

¹⁴The equivalent conditional logit (fixed-effect logit) estimate equals -0.523 and is insignificantly different from zero (p -value = 0.226). Note that this coefficient estimate can not be given a marginal effect interpretation. See Column (2) in Appendix Table A1.

Table 4: Estimated Effect of Price-Vector (6,8) on Buyer’s Interaction Decision

	Endogenous Price (<i>endo</i>)			Exogenous Price (<i>exo</i>)	
	(1)	(2)	(3)	(4)	(5)
vector (6,8)	0.199*** (0.040)	0.114** (0.053)	0.137*** (0.036)	-0.009 (0.037)	0.012 (0.031)
verifiability	0.102* (0.052)	0.118* (0.062)	-	0.142*** (0.052)	-
period	-0.030*** (0.003)	-0.032*** (0.003)	-0.030*** (0.003)	-0.025*** (0.003)	-0.026*** (0.003)
price p^l	-0.062*** (0.012)	-0.052*** (0.019)	-0.056*** (0.011)	-0.058*** (0.015)	-0.054*** (0.016)
price p^h	-0.100*** (-0.012)	-0.059*** (0.021)	-0.105*** (-0.013)	-	-
sample size	1472	1109	1472	1536	1536

Note: Figures in columns (1), (2) and (4) are coefficients from OLS regressions and figures in columns (3) and (5) are coefficients from customer fixed-effect regressions. Clustered standard errors shown in parentheses. *, ** and *** denote significance at .10, .05 and .01 levels.

we instead use a probit model, we find a marginal effect estimate of 23.2 percentage points (See Column (4) in Appendix Table A1). In column (2) we re-estimate the OLS regression model using only those observations where $p^l \leq 6$ and $p^h \leq 8$. In other words, we only compare price-vector (6,8) to vectors where both prices are lower. Even for this sub-sample we find a significant positive effect – customers are 11.4 percentage points more likely to accept the vector (6,8) than ”cheaper” vectors! Customer fixed-effect estimates, shown in column (3), also suggest that customers are more likely to interact with sellers posting price-vector (6,8).¹⁵

To determine whether (6,8) is the only price-vector for which there is a significantly positive interaction effect, we repeat our strategy from the test of H1 and re-estimate equation (3) with alternative price-vector dummy variables. We find that price-vector (3,5) also induces significantly more interaction (than would be predicted by linear p^l and p^h effects). However, this price-vector is chosen in only 1% of cases (15 out of 1472), compared to the 40% in which (6,8) is chosen. Furthermore, it is only significant at the 5% level ($p = 0.048$), and given the multiple hypothesis tests conducted in the sequential testing procedure and the consequent increase in type I error that it induces, it would be unwise to place too much stock in its significance. No other price-vector dummy variable is significantly positive at the 10% level.

To test the second part of H2, we re-estimate equation (3) using data from the *exo* treatments. With exogenous prices, the positive interaction effect of (6,8) should disappear if the relationship is driven solely by customers associating price-vector (6,8) with *g*-type sellers.

¹⁵Given that sellers and customers are randomly and anonymously matched, it is unlikely that the customer fixed-effect (represented by θ_i) would be correlated with prices - only by accident could customers with high propensities to interact be more or less likely to be confronted with (6,8). The similarity of estimates in columns (1) and (2) attest to this reasoning.

And indeed, the estimated (6,8) effects in columns (4) and (5) of Table 4, corresponding to OLS and customer fixed-effect regressions, show that in the *exo* treatments price-vector (6,8) does not induce more interaction.

4.3 Hypothesis 3 (H3)

The first part of Hypothesis 3 (H3) states that with respect to behavior in the price-posting stage of the game, two "varieties" of sellers should be evident, one that regularly posts (6,8), and another that rarely posts this price-vector. Table 2 shows that (6,8) is chosen more frequently than any other price-vector, indicating that (6,8) is in some sense special; but is it true that one group of sellers consistently chooses (6,8) while another group rarely does so, or are all sellers choosing (6,8) with roughly the same frequency? Figure 2 contains a histogram of the number of occasions each seller posts (6,8) in rounds 7 to 16 (i.e. the last 10 rounds). It shows that 47% of sellers never post (6,8), while 14% of sellers post (6,8) on every occasion. Furthermore, the never-posters and the always-posters are the only two mass-points in the data, with the remaining sellers posting (6,8) on one to nine occasions in roughly equal proportions.

Figure 2 also contains a dashed line graph overlaying the histogram. This graph is the hypothetical probability distribution that occurs if all sellers have an equal probability of 0.32 to post price-vector (6,8) in each of the last 10 rounds; the value 0.32 equals the proportion of times (6,8) is observed in the last 10 rounds (while 0.30 is the value for all 16 rounds). Clearly, this hypothetical behavior does not explain the data; however, to more formally test for the existence of a never-post group and an always-post group, we use a bootstrap procedure to test the null hypothesis that all sellers have an equal probability of posting price-vector (6,8). More specifically, we randomly draw samples of 92 sellers from a binomial distribution with probability p and $n = 10$. For example, we draw from the hypothetical probability distribution plotted in Figure 2. Our test statistic is the difference between the 90th and 10th percentiles, which equals 10 in our data (90th percentile = 10 times (6,8) posted; 10th percentile = 0 times (6,8) posted). We repeat the test for 99 hypothesized probabilities (0.01, 0.02, . . . , 0.98, 0.99), and on each occasion we use 1000 repetitions. The results from the testing procedure are definitive. In no repetition, for any of the 99 tested probabilities, does the randomly drawn data produce a 90-10 split equal to 10 (the maximum simulated value equalled 6). In other words, we can reject with extreme confidence the null hypothesis of a single seller type.

The second part of H3 states that sellers posting price-vector (6,8) are of two types, namely the selfish *b*-types and the good *g*-types. Figure 3 shows the propensity of sellers to provide appropriate quality when a customer's problem is major. This propensity is shown for sellers who rarely post price-vector (6,8) and for sellers who often post this vector.¹⁶ Roughly 40% of the often-posters always provide high quality when needed and roughly 40% never provide high quality when needed. This supports our theory with respect to the existence of two types of sellers posting the signalling price-vector, a 'good' type, which provides high quality when it is needed, and a 'bad' type, which always provides low quality. Amongst

¹⁶In the group of 'rarely-posters' we include all sellers who post (6,8) on two or fewer occasions in the last 10 rounds. The group of 'often-posters' includes all sellers who post (6,8) on eight or more occasions in the last 10 rounds.




Figure 2: Distribution of the Frequency Price-Vector (6,8) is Posted in the Last Ten Rounds

the rarely-posters, roughly 75% never provide high quality when needed, and roughly 20% always provide high quality when needed (conditional on customer interaction). The latter suggests that some good types are unable to ‘find’ the signalling price-vector.

5 Conclusion

We have proposed a simple model of a market for credence goods populated by three types of experts. ”The good” receive an extra utility from solving the consumer’s problem – for instance, because they value efficiency, or because they care for the well-being of consumers; ”the bad” are rational own-money maximizing experts – they take advantage of customers whenever there are material incentives for doing so; and ”the naive” are like ”the bad” in terms of preferences, but fall victim to a projection bias – they cannot conceive others being non-selfish. In this model we proved the existence of a semi-separating signalling equilibrium where experts posting ”fair” prices are more likely to provide efficient service than those posting different prices. The model implies a causal effect from other regarding preferences towards pricing behavior but no causal impact of prices on provision behavior. Specifically, in the proposed equilibrium ”the good” choose fair prices and behave in a consumer-friendly manner; ”the bad” mimic the good sellers in their price-setting behavior but take advantage of customers in their provision and charging decisions; and ”the naive” randomize in their price-setting behavior and take advantage of consumers if they find a victim. Customers understand the signalling value of prices and are hence more willing to trade if ”fair” prices

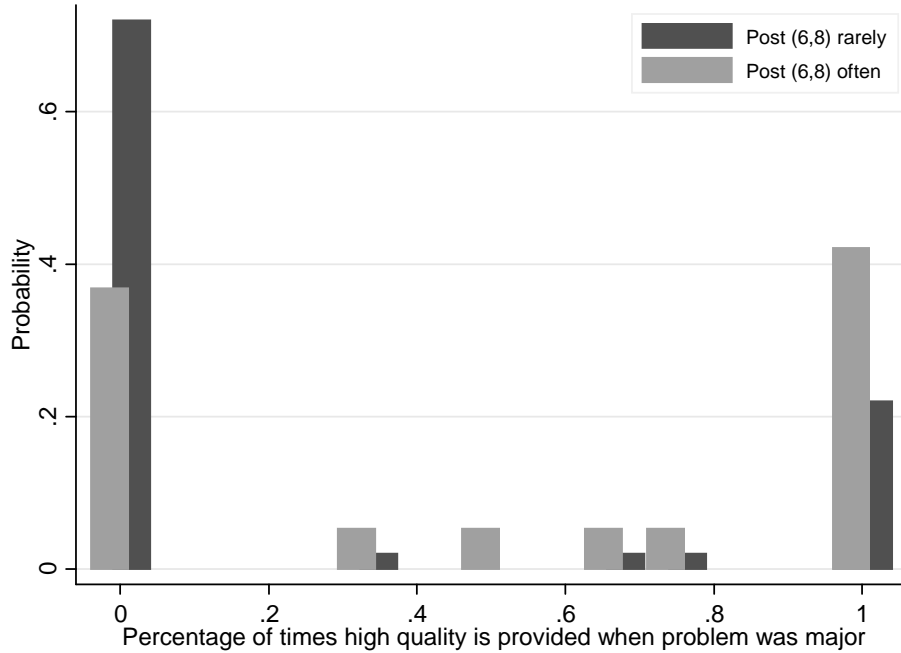


Figure 3: Distribution of the Probability that High Quality Is Provided when Problem is Major in the Last Ten Rounds

are posted.

Three hypotheses follow from this model: First, sellers who post the signalling price-vector are more prone to provide efficient service than those sellers posting a different price-vector. Second, customers are more willing to interact when the signalling price-vector is posted than when a different price-vector is posted. And third, sellers either consistently post the signalling price-vector or randomize over several price-vectors.

Using experimental data from sessions with endogenous and exogenous price choice we were able to provide strong evidence in support for each of the three hypotheses. In particular, fixed effect regressions as well as data from treatments with exogenous prices provided strong evidence indicating that efficient behavior is dependent on seller type, and that seller type determines prices. Therefore, posted prices do contain valuable information for consumers about the type of seller they face. Importantly, we have shown that customers do indeed recognize the signalling value of prices and respond to this information. Specifically, they are significantly more willing to accept if the signalling price-vector is posted than if any other price-vector is observed, even if the alternative vector has cheaper prices! Finally, the data has supported the hypothesis that there are three types of sellers, one (the naive one) that randomizes over the available price-vectors and two that post the signalling price-vector with high frequency, of the latter two types one (the good one) provides appropriate quality very frequently, the other (the bad one) behaves in line with the standard economic model of own-money maximizing behavior.

What do our findings imply for regulation of markets for credence goods? Standard

theory argues that if experts do not choose prices that align their incentives with customers' interest, regulation can help by externally imposing the correct pricing scheme. The results presented in this paper suggest that price-regulation might be counterproductive because it robs consumers of valuable information about the type of expert they face. This implies that endogenous price setting is a factor contributing to efficiency in credence goods markets.

References

- [1] Akerlof, George. 1970. "The Market for Lemons: Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics* 85(3): 488-500.
- [2] Andreoni, James and B. Douglas Bernheim. 2009. "Social Image and the 50-50 Norm: A Theoretical and Experimental Analysis of Audience Effects." *Econometrica* 77, 1607–1636.
- [3] Andreoni, James and John H. Miller. 2002. "Giving According to GARP: An Experimental Test of the Consistency of Preferences for Altruism." *Econometrica* 70, 737-53.
- [4] Angrist, Joshua D. and Jorn-Steffen Pischke. 2009. *Mostly Harmless Econometrics: An Empiricists Companion*. Princeton: Princeton University Press.
- [5] Balafoutas, Loukas, Adrian Beck, Rudolf Kerschbamer and Matthias Sutter. 2011. "What Drives Taxi Drivers? A Field Experiment on Fraud in a Market for Credence Goods", CESifo Working Paper 3461.
- [6] Berg, Joyce, John Dickhaut, and Kevin McCabe. 1995. "Trust, Reciprocity, and Social History." *Games and Economic Behavior* 10(1): 122-142.
- [7] Bolton, Gary E. and Axel Ockenfels. 2000. "ERC – A Theory of Equity, Reciprocity and Competition." *American Economic Review* 90(1): 166-193.
- [8] Charness, Gary and Brit Grosskopf. 2001. "Relative Payoffs and Happiness: An Experimental Study". *Journal of Economic Behavior and Organization* .45, 301-328.
- [9] Charness, Gary and Matthew Rabin. 2002. "Understanding Social Preferences with Simple Tests." *Quarterly Journal of Economics* 117(3): 817-869.
- [10] Darby, Michael R., and Edi Karni. 1973. "Free Competition and the Optimal Amount of Fraud." *Journal of Law and Economics* 16(1): 67-88.
- [11] Dulleck, Uwe, and Rudolf Kerschbamer. 2006. "On Doctors, Mechanics, and Computer Specialists: The Economics of Credence Goods." *Journal of Economic Literature* 44(1): 5-42.
- [12] Dulleck, Uwe, Rudolf Kerschbamer, and Matthias Sutter. 2011. "The Economics of Credence Goods: On the Role of Liability, Verifiability, Reputation and Competition." *American Economic Review* 101(2): 526-555.

- [13] Fehr, Ernst, and Klaus Schmidt. 1999. "A Theory of Fairness, Competition, and Cooperation." *Quarterly Journal of Economics* 114(3): 817-868.
- [14] Fischbacher, Urs. 2007. "Z-tree: Zürich Toolbox for Readymade Economic Experiments." *Experimental Economics* 10(2): 171-178.
- [15] Greiner, Ben. 2004. "An Online Recruitment System for Economic Experiments." In: Kremer, Kurt and Volker Macho (eds.): *Forschung und wissenschaftliches Rechnen 2003*. GWDG Bericht 63. Gesellschaft für Wissenschaftliche Datenverarbeitung, Göttingen, 79-93.
- [16] Gruber, Jonathan, John Kim, and Dina Mayzlin. 1999. "Physician Fees and Procedure Intensity: The Case of Cesarean Delivery." *Journal of Health Economics* 18(4): 473-490.
- [17] Gruber, Jonathan, and Maria Owings. 1996. "Physician Financial Incentives and Cesarean Section Delivery." *RAND Journal of Economics* 27(1): 99-123.
- [18] Holmes, David S. 1978. "Projection as a Defense Mechanism." *Psychological Bulletin* 85: 677-688.
- [19] Hubbard, Thomas N. 1998. "An Empirical Examination of Moral Hazard in the Vehicle Inspection Market." *RAND Journal of Economics* 29(2): 406-426.
- [20] Huck, Steffen, Gabriele Lünser, and Jean-Robert Tyran. 2006. "Competition Fosters Trust." UCL Working Paper.
- [21] Huck, Steffen, Gabriele Lünser, and Jean-Robert Tyran. 2007. "Pricing and Trust." UCL Working Paper.
- [22] Huck, Steffen, Gabriele Lünser, and Jean-Robert Tyran. 2010. "Consumer Networks and Firm Reputation: A First Experimental Investigation." *Economic Letters* 108(2): 242-244.
- [23] Iizuka, Toshiaki. 2007. "Experts' Agency Problems: Evidence From the Prescription Drug Market in Japan." *RAND Journal of Economics* 38(3): 844-862.
- [24] Iriberry, Nagore and Pedro Rey-Biel. 2010. Elicited Beliefs and Social Information in Modified Dictator Games: What do Dictator Believe Other Dictators Do? Economics Working Papers 1137, Department of Economics and Business, Universitat Pompeu Fabra.
- [25] Kerschbamer, Rudolf, Matthias Sutter, and Uwe Dulleck. 2009. "The Impact of Distributional Preferences on (Experimental) Markets for Expert Services." Working Paper, University of Innsbruck.
- [26] Loewenstein, George, Ted O'Donoghue and Matthew Rabin. 2003. "Projection Bias in Predicting Future Utility". *Quarterly Journal of Economics*, 118(4), 1209-1248.
- [27] McShane, Steven and Tony Travaglione. 2003. *Organisational Behaviour*, McGraw-Hill Australia.

- [28] Mullen, Brian, Jennifer L. Atkins, Debbie S. Champion, Cecelia Edwards, Dana Hardy, John E. Story and Mary Venderklok. 1985. "The False Consensus Effect: A Meta-Analysis of 115 Hypothesis Tests", *Journal of Experimental Social Psychology* 21, 263-283.
- [29] Ross, Lee, David Greene and Pamela House. 1977. "The 'False Consensus Effect': An Egocentric Bias in Social Perception and Attribution Processes". *Journal of Experimental Social Psychology* 13, 279-301.
- [30] Selten, Reinhard and Axel Ockenfels. 1993. "An Experimental Solidarity Game". *Journal of Economic Behavior and Organization* 34, 517-539.
- [31] Sherwood Gregory G. 1981. "Self-Serving Biases in Person Perception: A Re-examination of Projection as a Mechanism of Defense." *Psychological Bulletin* 90, 445-459.
- [32] Wolinsky, Asher. 1993. "Competition in a Market for Informed Experts' Services." *RAND Journal of Economics* 24(3): 380-398.
- [33] Wolinsky, Asher. 1995. "Competition in Markets for Credence Goods." *Journal of Institutional and Theoretical Economics* 151(1): 117-131.

Appendix

Appendix A: Additional Regressions

Table A1: Estimated Effects From Binary Choice Models

	Seller's Provision Decision			Buyer's Interaction Decision	
	(1)	(2)	(3)	(4)	(5)
vector (6,8)	0.121*** (0.046)	-0.523 (0.432)	-0.038 (0.038)	0.232*** (0.042)	-0.009 (0.039)
verifiability	-0.081 (0.055)			0.125** (0.059)	0.154*** (0.056)
period	0.001 (0.004)	-0.026 (0.033)	-0.001 (0.004)	-0.035*** (0.004)	-0.027*** (0.003)
major problem	-0.526*** (0.049)	-3.518*** (0.312)	-0.259*** (0.041)		
price p^l	0.005 (0.018)	-0.003 (0.170)	0.028** (0.013)	-0.077*** (0.015)	-0.062*** (0.016)
price p^h	0.056** (0.023)	0.435** (0.211)		-0.126*** (0.016)	
sample size	697	697	903	1472	1536

Note: Columns (1), (3), (4) and (5) present estimated probit marginal effects evaluated at mean values, and correspond to Column (2) Table 3, Column (6) Table 3, Column (1) Table 4, and Column (4) Table 4. Column (2) presents coefficient estimates from a fixed-effect logit, and corresponds to column (4) Table 3.

Appendix B: Experimental Instructions

(for treatment *NV-endo* – instructions for the other treatments are available upon request)

INSTRUCTIONS FOR THE EXPERIMENT

Thank you for participating in this experiment. Please do not talk to any other participant until the experiment is over.

2 Roles and 16 Rounds

This experiment consists of **16 rounds**, each of which consists of the same sequence of decisions. This sequence of decisions is explained in detail below. There are 2 kinds of roles in this experiment: **player A** and **player B**. At the beginning of the experiment you will be randomly assigned to one of these two roles. On the first screen of the experiment you will see which role you are assigned to. Your role remains the same throughout the experiment. A player A interacts with a player B. This pair of players **changes** for each round. Therefore you are interacting in every round with a **new** player (of the other role).

All participants get the same information on the rules of the game, including the costs and payoffs of both players.

Overview of the Sequence of Decisions in a Round

Each round consists of a maximum of 4 decisions which are made consecutively. Decisions 1, 3 and 4 are made by player A, decision 2 is made by player B.

Short Overview of the Sequence of Decisions in a Round

1. Player A chooses one price for action 1 and one price for action 2.
2. Player B gets to know the prices chosen by the player A. Then player B decides whether he/she wants to interact with player A. If not, this round ends.

If player B chooses to interact then

3. Player A (but **not** player B) is informed about the type of player B. There are two possible types of player B: he/she is of either type 1 or type 2. Player A has to choose an action: either action 1 or action 2.

4. Player A then decides between one of the two prices chosen in decision 1. This price has to be paid by Player B. The price chosen does not need to be the price of the action chosen in decision 3, it can be other price.

Detailed Illustration of the Decisions and Their Consequences Regarding Payoffs

Decision 1

In case of an interaction each **player A** has to choose between 2 actions (action 1 and action 2) at decision 3. Each chosen action causes costs which are as follows:

Action 1 costs player A 2 points (= currency of the experiment).

Action 2 costs player A 6 points.

Player A can charge prices for these actions from all those players B who decide to interact with him/her. At **decision 1** each Player A has to set the **prices for both actions**. Only (strictly) positive integer numbers are possible, i.e., only 1, 2, 3, 4, 5, 6, 7, 8, 9, 10 and 11 are valid prices. Note that the price for action 1 must not exceed the price for action 2.

Decision 2

Player B gets to know the prices set by player A for the two actions at decision 1. Then player B decides whether he/she wants to interact player A.

If he/she wants to do so, player A can choose an action at decision 3 and charge a price in decision 4 (see below). Player B will not be able to observe which action player A chooses.

If he/she doesn't want to interact, this round **ends** and both players get a **payoff of 1.6 points for this round**.

Decision 3

Before decision 3 is made (in case player B chose "Yes" at decision 2) a type is randomly assigned to player B. **Player B** can be of one of two types: **type 1** or **type 2**. This type is determined **new** for each player B in **each round**. With a **probability of 50%** player B is of **type 1**, and with a **probability of 50%** he/she is of **type 2**. Imagine that a coin is tossed for each player B in each round. If the result is e.g. "heads", player B is of type 1, if the result is "tails" he/she is of type 2.

Every player A gets to know types player B before he/she makes decision 3. Then player A chooses an action, either action 1 or action 2.

An action is sufficient under the following conditions:

- a) In case player B is a type 1 player and player A can chose either action 1 or action 2.
- b) In case player B is a type 2 player and player A chooses action 2.

An action is not sufficient if player B is of type 2 and player A chooses action 1.

Player B receives 10 points, if the **action** chosen by player A is **sufficient**. **Player B** receives 0 points if the **action** chosen by player A is **not sufficient**.

At no time player B will be informed whether he/she is of type 1 or a type 2 player or which action player A chose.

Decision 4

Player A **charges** one of the two **prices**, determined in decision 1, for the action he chose in decision 3. The price charged does not have to be the price of the action chosen in decision 3, it can be the price of the other action.

Payoffs

If player B chose not to interact in decision 2 (*decision “No” of player B*) then both players receive **1.6 points** for this round.

Otherwise (*decision “Yes” by player B*) the payoffs are as follows:

Player A receives the **price** (denoted in points) charged in decision 4 **less the costs** for the action chosen in decision 3.

Player B's payoff depends on whether the decisions 3 of player A was sufficient or not.

- a) The action was sufficient. **Player B** gets **10 points less** the price charged at decision 4.
- b) The action was not sufficient. **Player B** has to pay the price charged at decision 4.

At the beginning of the experiment you receive an initial **endowment of 6 points**. With this endowment you are able to cover losses that might occur in some rounds. Losses can also be compensated by gains in other rounds. If your total payoff sums up to a loss at the end of the experiment you will have to pay this amount to the supervisor of the experiment. By participating in this experiment you agree to this term. Please note that there is **always** a possibility to avoid losses in this experiment.

To calculate the final payoff the initial endowment and the profits of all rounds are added up. This sum is then converted into cash using the following exchange rate:

1 point = 25 Euro-cents
(i.e. 4 points = 1 Euro)