

IZA DP No. 4647

The Impact of Distributional Preferences on (Experimental) Markets for Expert Services

Rudolf Kerschbamer
Matthias Sutter
Uwe Dulleck

December 2009

The Impact of Distributional Preferences on (Experimental) Markets for Expert Services

Rudolf Kerschbamer

University of Innsbruck

Matthias Sutter

*University of Innsbruck,
University of Gothenburg and IZA*

Uwe Dulleck

Queensland University of Technology

Discussion Paper No. 4647

December 2009

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 Impact of Distributional Preferences on (Experimental) Markets for Expert Services

Credence goods markets suffer from inefficiencies arising from informational asymmetries between expert sellers and customers. While standard theory predicts that inefficiencies disappear if customers can verify the quality received, verifiability fails to yield efficiency in experiments with endogenous prices. We identify heterogeneous distributional preferences as the main cause and design a parsimonious experiment with exogenous prices that allows classifying experts as either selfish, efficiency loving, inequality averse, inequality loving or competitive. Results show that most subjects exhibit non-standard distributional preferences, among which efficiency-loving and inequality aversion are most frequent. We discuss implications for institutional design and agent selection in credence goods markets.

JEL Classification: C72, C91, D82

Keywords: distributional preferences, credence goods, verifiability, experiment

Corresponding author:

Matthias Sutter
University of Innsbruck
Department of Public Finance
Universitaetsstrasse 15
A-6020 Innsbruck
Austria
E-mail: matthias.sutter@uibk.ac.at

1 Introduction

Credence goods markets suffer from informational asymmetries between expert sellers and customers, because customers are unable to identify the quality they need, whereas expert sellers are able to do so. Darby and Karni (1973) added credence goods to Nelson's (1970) classification of ordinary, search and experience goods, and mention provision of repair services as typical example. Other important examples of credence goods are health care provision, financial consulting, or sales services of technical equipment.

The informational asymmetries on credence goods markets may cause a variety of inefficiencies. Depending on informational conditions and prices for different qualities, expert sellers may have a material self interest to provide unnecessarily high quality (a case referred to as "overtreatment"), or insufficiently low quality ("undertreatment"), or to charge for a higher quality than provided ("overcharging"). Such inefficiencies are not only a theoretical possibility, but are well documented in the literature. A survey conducted by the Department of Transportation indicates that more than half of car repairs are unnecessary, which is an indication of overtreatment (Wolinsky 1993, 1995). Schneider (2006) presents a field study on fraudulent behavior in the auto repair business. He finds substantial evidence of all three kinds of fraud, (i) overtreatment (in the form of completely unnecessary repairs), (ii) undertreatment (by neglecting defects that require urgent attention), and (iii) overcharging (billing for parts and labor not provided). Referring to the medical sector, Gruber and Owings (1996) or Gruber, Kim and Mayzlin (1999) show that the relative frequency of Cesarean deliveries compared to normal child births depends on the fee differentials of health insurance programs for both types of treatments, despite these differentials being unrelated to medical indication. Hence, monetary incentives matter for the provision of health care and for the type of inefficiency arising in market equilibrium.¹

A common finding in the theoretical literature is that introducing verifiability ensures efficiency on credence goods markets. Verifiability is defined to apply when an expert seller cannot charge for a treatment quality that has not been provided.² If verifiability applies, experts are predicted to choose

¹Experimental evidence on inefficiencies in credence goods markets is presented in Dulleck, Kerschbamer and Sutter (2009), while the related studies by Huck, Lünser and Tyran (2006, 2007) focus on experience goods (goods –like wine– where the quality cannot be observed ex ante but is perfectly observable ex post, i.e. after consumption).

²Verifiability of the provided treatment quality is likely to hold in many important

prices for different quality levels that induce them to provide the appropriate quality of the credence good. As a consequence, consumers - inferring experts' incentives from posted prices - decide to interact.³

In this paper we present the results of laboratory experiments with endogenous price choices by experts showing that - contrary to the theoretical predictions - verifiability is of little help in promoting efficiency on credence goods markets. In fact, we find that the relative frequencies of market interaction, undertreatment and overtreatment do not differ significantly between two experimental conditions that are identical except that verifiability applies in one condition, but not the other.⁴ The aggregate performance in both conditions is better than the standard prediction for a market without verifiability, but worse than the standard prediction for a market with verifiability. These findings raise two important research questions: Why is the performance on credence goods markets so poor even in the presence of verifiability? And why is the performance of credence goods markets without institutional safeguards against fraud (i.e. without verifiability) so much better than predicted? While this paper's main focus is on the former question, we will explain below that our results also help to answer the latter one.

We start by presenting evidence indicating that non-standard distributional preferences play an important rule for the behavior of market par-

credence goods markets. In some cases, for example pest control, equipment repair and dental services, the customer is present during treatment and can ensure that services charged for are actually provided. In other cases verifiability is secured indirectly through the provision of ex post evidence. For instance, before stricter hygienic rules were enacted in hospitals, it was not uncommon for surgeons performing gall stone extractions to provide the extracted stones to the patient after the operation to show that the suggested intervention had actually been performed (and to prove that it was necessary). Similarly, in today's automobile repair market, it is quite common that broken parts are handed over to the customer to substantiate the claim (on the bill) that replacement, and not only repair, has been performed.

³Theoretical studies investigating credence goods markets under verifiability include Emons (1997 and 2001), Alger and Salanié (2002), Pesendorfer and Wolinsky (2003), and Dulleck and Kerschbamer (2008 and 2009). See Dulleck and Kerschbamer (2006) for a unifying model and a survey of the literature.

⁴These two experimental conditions are also included in a companion paper (Dulleck et al. 2009) where we investigate with a total of 16 experimental treatments the role of liability, verifiability, competition and reputation on the frequency of trade and the efficiency on credence goods markets. In the companion paper we do not address the issue of how distributional preferences affect provision behavior, nor do we present a test for identifying distributional preferences of experts (as we do here).

ticipants in our endogenous price choice experiments. Motivated by these findings we proceed with a thorough theoretical analysis of the impact of distributional preferences on the provision behavior of experts under verifiability. Based on our theoretical results that require minimal assumptions and do not rely on any structural assumption on experts' utility or motivation function (such as linearity, or convexity or whatsoever), we then design a parsimonious experiment allowing us to classify experimental sellers' distributional preferences as either selfish, inequality averse, efficiency loving, inequality loving or competitive. Our experimental design relies on exogenously given prices for different quality levels of a credence good rather than letting sellers decide endogenously as in the first set of experiments. The main advantage of using a fixed price design is that it enables us to elicit an expert's willingness to pay for the implementation of her preferred payoff distribution while avoiding endogeneity problems of certain types choosing particular price vectors.

We implement our experiment and find that the behavior of a large majority of subjects is consistent with either a taste for efficiency (in the spirit of Charness and Rabin 2002) or inequality aversion (in the spirit of Fehr and Schmidt 1999, or Bolton and Ockenfels 2000). Only about 20% of subjects behave as standard theory predicts.

How do these findings relate to the above mentioned research questions? The key insight is that one and the same heterogeneity in distributional preferences can lead to a positive or a negative deviation from standard theory's prediction depending on institutional design and - within a given design - on the price vector under which the transaction takes place. We will argue that without verifiability the negative side of distributional preferences cannot lead to a deviation from standard theory's prediction while the positive side can and does. With verifiability exactly the opposite is true. In sum, the net effect of distributional preferences is positive in markets without verifiability, but negative with verifiability, leading to both markets performing equally (well or poor) in the aggregate.

More generally speaking, our results show that distributional preferences can have both considerable up- and downside effects on efficiency in markets with asymmetric information. This confirms earlier findings of Fehr, Kirchsteiger and Riedl (1993) and Fehr, Gaechter and Kirchsteiger (1997). Fehr et al. (1993) document that fairness preferences can prevent market clearing in labour markets, while the results in Fehr et al. (1997) show that social preferences can increase - rather than decrease - efficiency in a labour

market with incomplete contracts. Our results add to the existing literature the important insight that institutions need to be robust against the dark side of distributional preferences of market participants in order to achieve efficiency. In credence goods markets verifiability does not meet this requirement despite standard theory's predictions.

In sum, our paper contributes to the existing literature in four dimensions. First it presents a thorough theoretical analysis of the impact of distributional preferences on expert provision behavior under verifiability. Secondly, it develops a parsimonious experiment that allows classifying experimental experts as either selfish, inequality averse, efficiency loving, inequality loving or competitive using minimal assumptions. Third, it applies the parsimonious experimental design (with exogenous prices) in the lab and shows that a substantial part of the subject population exhibits non-standard preferences. And fourth, it discusses the implications of heterogeneity in distributional preferences for institutional design and agent selection in credence goods markets.

The remainder of the paper is organized as follows. Section 2 introduces the basic setup of a credence goods market and then presents standard predictions and the experimental results of two treatments with endogenous prices which only differ with respect to the presence or absence of verifiability. There we also discuss evidence indicating that distributional preferences are important determinants of market participants' behavior. In Section 3 we, first, investigate theoretically the behavioral implications of two competing models of other-regarding preferences, one based on preferences for efficiency and one based on inequality aversion. Second, we extend our focus to a more general discussion on the impact of different kinds of distributional preferences on provision and charging behavior in credence goods markets. This discussion culminates in the design of the parsimonious experiment which is then implemented in Section 4 to classify expert sellers in different distributional types according to their provision behavior. Section 5 concludes with a discussion of our results and their implications for institutional design (institutions should be robust against the coexistence of different types of agents) and agent selection (selecting the right agents converts verifiability into a robust institution, in the sense that it performs well even under cost uncertainty - a case largely ignored in the theoretical literature).

2 The Role of Verifiability in Credence Goods Markets: A Simple Model, Standard Theory's Predictions and Experimental Evidence

This section introduces a simple model of a credence goods market and shows how verifiability should affect behavior according to predictions based on standard preferences (i.e., assuming that rational subjects maximize their own material payoffs). It then presents experimental evidence showing that verifiability does not have the predicted effects and indicating that non-standard distributional preferences shape the behavior of market participants.

2.1 Basic Model

Consumers are *ex ante* identical. Using a medical farming, we assume that customers need a major treatment (t^h) with probability h , and a minor treatment (t^l) with probability $1 - h$. Each consumer (he) is randomly matched with one seller (she) who sets prices p^h and p^l for the major, respectively minor, treatment (with $p^h \geq p^l$). The seller has costs c^h for the major treatment, and c^l for the minor one (with $c^h > c^l$).

The consumer only knows the prices for the different treatments, but not the type of treatment that he needs, when he makes his decision whether or not to interact with the seller. In case of interaction, the seller gets to know which type of treatment the customer needs. Then she provides one of the two treatments and charges one of the two prices.

Customers in need of the minor treatment t^l are sufficiently treated in any case (both if the seller chooses t^l and if she chooses t^h). However, if the customer needs the major treatment t^h , then only t^h is sufficient. A sufficient treatment yields a value $v > 0$ for the customer, an insufficient treatment yields a value of zero. If the customer decides against interaction then both the customer and the seller receive an outside option of $o \geq 0$. In case of an interaction, the monetary payoff for the consumer is the value from being treated minus the price to be paid, whereas the seller receives as a monetary payoff the price charged minus the costs of the treatment provided. More formally, let $\theta \in \{l, h\}$ be the index of a customer's type of problem, $\mu \in \{l, h\}$ the index of the quality of treatment provided, and $\kappa \in \{l, h\}$ the index of the quality of treatment charged for. Then the material payoff of




Figure 1: The Credence Goods Game

the seller in case of interaction is

$$\pi_s(p^l, p^h, \mu, \kappa) = p^\kappa - c^\mu \quad (1)$$

while the consumer's material payoff in case of interaction is

$$\pi_c(p^l, p^h, \theta, \mu, \kappa) = v - p^\kappa \text{ if } \theta \leq \mu \text{ and } -p^\kappa \text{ otherwise.}^5 \quad (2)$$

Figure 1 presents this game. Note, this simple game captures all the idiosyncratic problems of credence goods markets discussed earlier. If a customer needs high quality, i.e. on the third stage nature moves left, and the seller treats with t^l , we have a situation of *undertreatment*, if he needs low quality and the seller treats with t^h we have *overtreatment*; and if the seller charges p^h when treatment t^l was provided, we have a situation of *overcharging*.

⁵Here we use the convention that $l \leq h$, but not vice versa.

2.2 Experimental Design

To understand the role of verifiability, we compare two experimental conditions.⁶ In one - condition N - we impose no verifiability and in the other - condition V - we impose verifiability. Condition N is exactly as described above and depicted in Figure 1. Verifiability means that consumers are able to observe and verify *ex post* the treatment that has been provided by the seller (without knowing, however, whether this treatment is the appropriate one). Therefore, condition V is identical to condition N, except that the last stage is degenerate because the expert has to charge the price for the provided treatment. This means verifiability prevents overcharging (and undercharging), but it does not preclude over- and/or undertreatment.

In both experimental conditions we let the customer's probability of needing the major treatment be $h = 0.5$, and the value of a sufficient treatment be $v = 10$. The costs of providing the minor, respectively major, treatment are $c^l = 2$, and $c^h = 6$. The prices posted by the sellers, p^l and p^h (with $p^l \leq p^h$), have to be chosen in integer numbers from the interval $\{1, 11\}$. The outside option if no trade takes place between the seller and the customer is set to $o = 1.6$.

We always use matching groups of eight subjects each, which is common knowledge in all conditions. Four subjects in each matching group are in the role of customers, and four in the role of sellers. The assignment to roles is randomly determined at the beginning of the experiment, and roles are kept fixed throughout the entire experiment. There are 16 periods of interaction in which we use a stranger matching with random rematching of experts and customers after each period to avoid an opportunity for reputation building.⁷

All experimental sessions were run computerized using zTree (Fischbacher 2007) and recruiting was done via ORSEE (Greiner 2004). A total of 184 subjects participated in this set of experiments, most of them studying economics, business administration, and social and life sciences. All sessions started with an extensive description of the game. All parameters as well as the matching procedure were made common knowledge to all participants by reading them aloud. Before the experiment started, participants had to

⁶Since in our model we let sellers provide a "treatment" to customers we refer to experimental treatments as conditions throughout the article.

⁷Of course, the probability of meeting a particular expert again is one quarter, meaning that implicit reputation formation that encompasses the whole matching group is feasible. However, the stranger matching precludes reputation formation of individual experts.

answer a set of control questions correctly to ensure that they had fully understood the instructions. The average session length was 1.5 hours, and subjects earned on average 15 Euro.

2.3 Standard Predictions and the Role of Verifiability

Standard theory (relying on rational, own-money-maximizing, risk-neutral agents) predicts that without verifiability the expert will always charge the higher price p^h and always provide the cheaper treatment t^l . Anticipating this consumers will only accept if $p^h \leq (1 - h)v - o = 3.4$. But with such a p^h even a defrauding expert earns less than the value of her outside option (because $(1 - h)v - c^l < 2o$). Thus, in condition N standard theory predicts that the market breaks down (on the equilibrium path). Furthermore, if - for whatever reason - a transaction takes place (off the equilibrium path), we should observe undertreatment and overcharging:

Prediction 1a (Condition N) *Consider Condition N. Assume that subjects have standard preferences. Then no interaction will take place (on the equilibrium path), and if an interaction takes place (off the equilibrium path) then the expert seller provides t^l and charges p^h .*

In condition V the expert cannot charge for a treatment other than the provided one and the provided treatment depends on the mark-up $p^i - c^i$, $i \in \{l, h\}$. To characterize the provision behavior of own-money-maximizing experts under this condition it is useful to distinguish the following types of price vectors:

- an equal mark-up price-vector is defined as one that satisfies $p^h - c^h = p^l - c^l$;
- an undertreatment price-vector satisfies $p^h - c^h < p^l - c^l$; and
- an overtreatment price-vector is characterized by $p^h - c^h > p^l - c^l$.

Under equal mark-up price vectors own-money-maximizing sellers provide the appropriate treatment⁸, and they provide always the minor (major)

⁸This is either by assumption, i.e. that if indifferent the expert will provide in the best interest of the customer, or by referring to the limit of a mixed strategy equilibrium. See Dulleck and Kerschbamer (2008) on this.

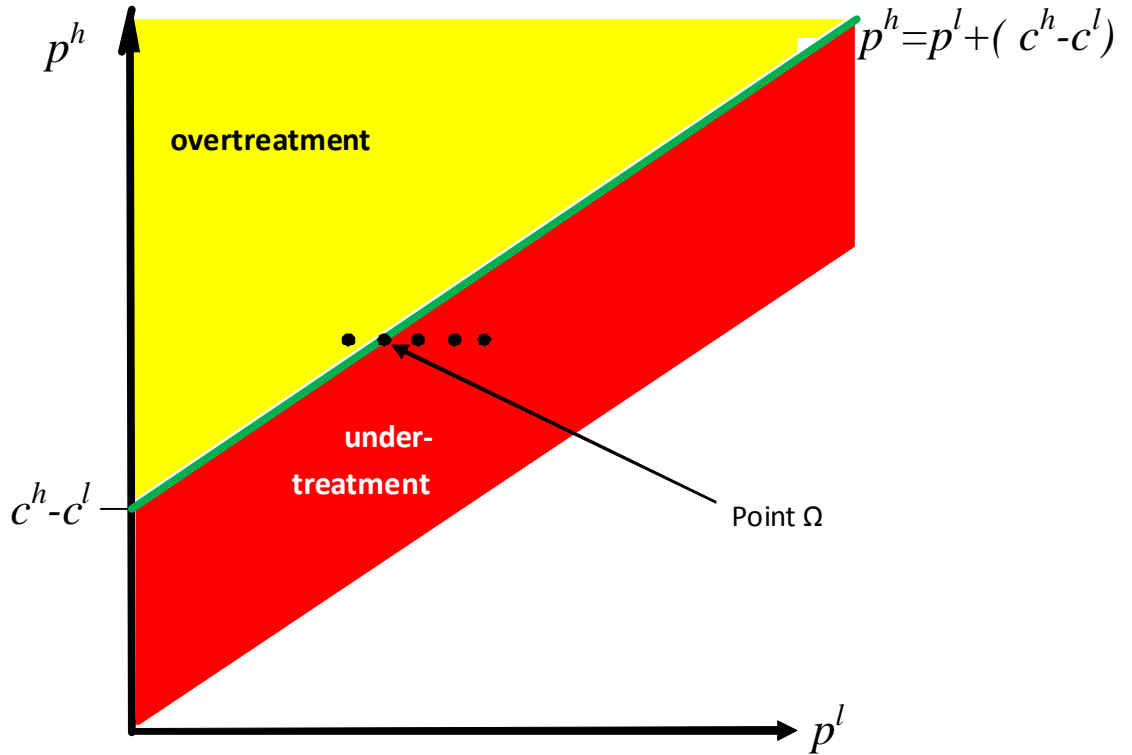


Figure 2: Provision Behavior of an Expert with Standard Preferences under Verifiability

treatment under undertreatment (overtreatment) price vectors. Figure 2 shows in the space of price vectors (p^h, p^l) the set of vectors that induce efficient service.⁹ Those price vectors lie on a line connecting all points where $p^h - c^h = p^l - c^l$. The line intersects the vertical axis at $p^h = c^h - c^l$ and has a slope of one. Below this line the expert is induced to always provide low quality (undertreatment in case the consumer needs t^h) and above the line the expert is induced to always provide high quality (overtreatment in case the consumer needs t^l).

Anticipating how experts' provision (and charging) behavior depends on price vectors, consumers will accept an equal mark-up vector iff $p^h \leq 10$, an

⁹"Point Ω " and of the five price vectors indicated by bullet points in Figure 2 are not important for the arguments in this section. We will refer to them in Subsection 3.3.

undertreatment vector iff $p^l \leq 3$, and an overtreatment vector iff $p^h \leq 8$. Thus, to maximize profits, experts will choose $p^h = 10$ and $p^l = 6$ which will be accepted by an own-money-maximising, risk-neutral consumer. This leads to the following prediction:

Prediction 1b (Condition V) *Assuming standard preferences, experts post $p^h = 10$ and $p^l = 6$ in condition V and consumers choose to enter the market and get appropriate treatment (on the equilibrium path). Under other price vectors customers get either always t^l (under undertreatment vectors), or always t^h (under overtreatment vectors), or always the appropriate treatment (under equal-mark-up vectors) depending on mark-ups (off the equilibrium path).*

We summarize our predictions 1a and 1b in:

Prediction 1 *Assume that subjects have standard preferences. Then in condition N no interaction will take place on the experimental credence goods market while in condition V all interactions will be carried out and full efficiency prevails.*

2.4 Experimental Results

2.4.1 Aggregate Behavior

Being aware that subjects in experiments often do not behave exactly as standard theory predicts, we expected to find a difference in the behavior in conditions N and V, but a less extreme one than in the standard benchmark. To our surprise we found no difference at all:

Observation 1 (Aggregate Behavior in Conditions N and V) *Verifiability has no significant impact on key variables in the aggregate: The frequency of interaction, the undertreatment rate, the overtreatment rate and overall efficiency are not significantly different in conditions V and N. The overall performance in both conditions is better than the standard prediction for condition N, but worse than the standard prediction for condition V.*

Table 1, Figure 3 and Figure 4 support this observation, leading us to reject Prediction 1. Predictions 1a and 1b also included predictions on off-equilibrium behavior. Recall that in condition N the undertreatment rate

should be (close to) one! This is obviously not what we observe: The undertreatment rate is high (53%) but far from the predicted 100%. Additionally, it is not significantly higher (and in fact, on average, lower) in condition N than in condition V.

Table 1: Summary Statistics for Conditions N and V

Averages per Period	Condition N	Condition V
Interaction	0.45	0.50
Efficiency ^a	0.18	0.16
Undertreatment ^b	0.53	0.60
Overtreatment ^c	0.06	0.05
Overcharging ^d	0.88	-
Profits Sellers	2.69	2.58
Profits Customers	1.00	1.06

None of the variables is significantly different between conditions N and V (using two-sided Mann-Whitney U-tests with matching groups of 8 subjects as one independent observation).

^a calculated as (actual average profit – outside option) divided by (maximal possible average profit – outside option)

^b customer needs t^h , but seller provides t^l

^c customer needs t^l , but seller provides t^h

^d seller provides t^l , but charges t^h (with $p^h > p^l$ and customer needs t^l)

What went wrong in condition V? According to the theoretical prediction sellers should choose equal mark-up prices. However, such prices are very rare in condition V - they are chosen in less than 5% of all transactions. Similarly rare are overtreatment price vectors, i.e. price vectors that provide material incentives to overtreat the customer (since $p^h - c^h > p^l - c^l$). Most posted price vectors are of the undertreatment type, i.e. $p^l - c^l > p^h - c^h$.

Table 2 below reports the frequencies of the five most popular price vectors posted by sellers in conditions N and V. It is interesting to note that in condition V only one equal mark-up vector, i.e. (4,8), is among the top 5 price vectors, but it is not the predicted one. In both conditions the price vector (6,8) is by far the most frequently posted price vector. This price vector splits the gains from trade equally between consumers and sellers - if sellers always provide the appropriate treatment and charge for the provided

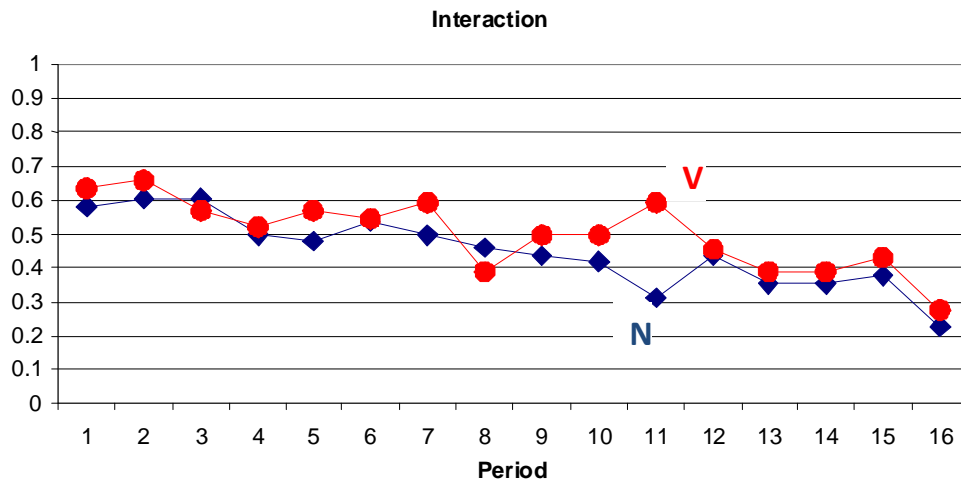


Figure 3: Relative Frequency of Interaction in Conditions N and V.

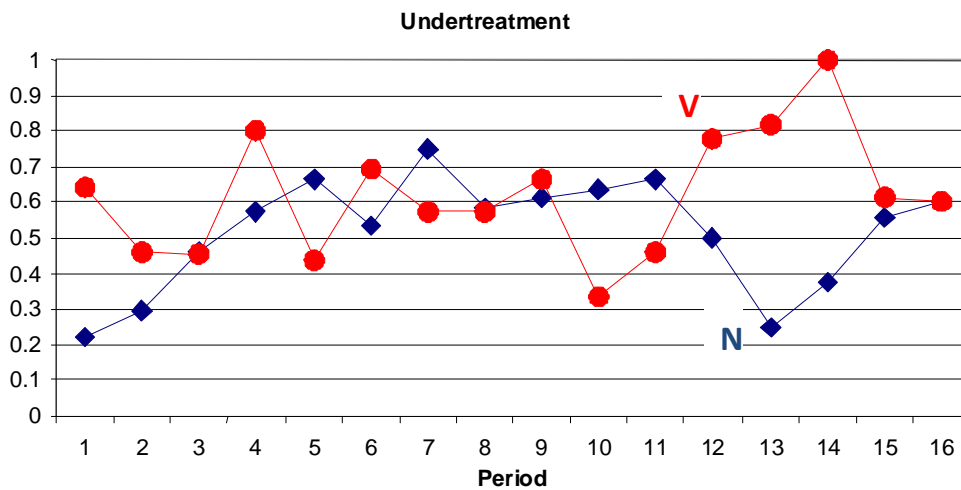


Figure 4: Relative Frequency of Undertreatment (Seller Provides t^l when t^h is Needed) in Conditions N and V.

treatment. The prominence of this price vector suggests that a concern for relative payoffs plays a role for aggregate behavior in the experiment.

Table 2: Most Popular Price Vectors in Conditions N and V

Condition N			Condition V		
(p^l, p^h)	absolute #	rel. frequency	(p^l, p^h)	absolute #	rel. frequency
(6,8)	176	22.92%	(6,8)	265	37.64%
(4,8)	84	10.94%	(7,8)	89	12.64%
(5,7)	50	6.51%	(5,8)	46	6.53%
(5,8)	44	5.73%	(4,8)	17	2.41%
(4,7)	39	5.08%	(8,8)	15	2.13%
393 (of 768)		51.17%	432 (of 704)		61.36%

Table 3 shows experts' provision behavior for the top 5 price vectors in condition V; more precisely, it shows how aggregate under- and overtreatment rates change in the price difference $p^h - p^l$. Note that the sign of the change in provision behavior is consistent with the direction predicted by standard theory (for off-the-equilibrium-path behavior): Increasing p^l while keeping p^h constant decreases the overtreatment rate and increases the undertreatment rate. However, there are no discontinuous jumps in provision behavior as predicted by standard theory. Also, the equal mark-up vector (4,8) induces a considerable amount of overtreatment instead of inducing appropriate treatment. As we will show later, this is exactly what theories of inequality aversion would predict independent of the functional form in which inequality averse preferences are modelled (let alone on specific parameterizations of a specific functional form). That is, we will show that *any* model of inequality aversion that has the equal split as the reference point *must* predict overtreatment under this price vector.

Table 3: Under- and Overtreatment Rates in Condition V

(p^l, p^h)	Overtreatment Rate	Undertreatment Rate
(4,8)	37.5%	0%
(5,8)	14.3%	33%
(6,8)	1.25%	53%
(7,8)	0%	65%
(8,8)	0%	100%

2.4.2 Individual Behavior

We also analyzed the behavior on the individual level.

Observation 2 (Individual Behavior in Conditions N and V) *In both conditions, N and V, there exist two types of experts that exhibit the same behavior throughout the 16 periods of the game. The first group consists of 'underproviders'. In condition N this group consists of 30% of the expert population (about 50% if only the behavior in the last ten periods is considered) and its members always provide low quality and always charge for high quality. In condition V this group consists of 40% of the subjects (again 50% if only the behavior in the last ten periods is considered) and its members always choose undertreatment price vectors and always provide low quality. The second group consists of 'appropriate providers'. In condition N this group consists of 25% of the expert population, in condition V it consists of 16% of the population. Members of this group always provide the appropriate quality in both conditions (although they post predominantly undertreatment vectors in condition V and although they are predicted to undertreat under each price vector in condition N).*

Table 4 illustrates this observation by showing the number of sellers being classified as those either underproviding all the time or consistently providing appropriate treatment. A possible explanation for appropriate treatment even when monetary incentives call for undertreatment is experts having a taste for efficiency. Another possible explanation is that experts care for an equitable payoff. Support for the latter hypothesis comes from the analysis of pricing behavior of different types of experts: In condition N, more than 50% of the (6,8) price vectors are posted by experts always providing the appropriate treatment while they make up only 25% of the population. In condition V more than 20% of the (6,8) price vectors are posted by experts always providing the appropriate treatment, while they make up only 16% of the population. Thus, those experts that tend to post the price vector that splits the gains from trade equally between consumers and sellers - if sellers always provide the appropriate treatment and charge for the provided treatment - are more likely to be appropriate-providers than others. Of course, this is only a very rough indication that distributional preferences shape behavior. As we will see later, we can say more in an experiment where price vectors are exogenously imposed.

Taken together, the evidence from the experiment with endogenously determined price vectors clearly shows both that verifiability does not yield full efficiency and that it also indicates that non-standard distributional preferences play a role for experts' provision and charging behavior.

Table 4: Number and Frequency of Subjects with Consistent Behavior over all 16 Periods

(in parenthesis values for the last ten periods)

type of behavior	Condition N		Condition V	
	absolute #	rel. frequency	absolute #	rel. frequency
underprovision	14 (25)	29(52)%	17 (21)	39(48)%
appropriate provision	12 (12)	25(25)%	7 (8)	16(18)%

3 Identifying Distributional Preferences and Their Impact on Provision Behavior

The main purpose of this section is to develop a parsimonious experimental design that allows classifying experimental experts according to their distributional preferences. In sections 3.1 and 3.2 we illustrate the impact of distributional preferences on experts' provision behavior by discussing the behavioral implications of the two most often invoked types of distributional preferences, i.e., efficiency concerns (e.g., Charness and Rabin 2002) and inequality aversion (e.g., Fehr and Schmidt 1999; Bolton and Ockenfels 2000). This illustration will be followed in Section 3.3 by the key contribution of this paper, i.e., a rigorous generalization of the impact of distributional preferences on the provision behavior on credence goods markets.

3.1 Preferences for Efficiency

Suppose that (some) sellers derive an extra utility from providing the appropriate treatment or that they feel a moral cost if they under- or overtreat a customer. To explore the consequences of such concerns for provision behavior let us start with a simple linear model. As before, let $\theta \in \{l, h\}$ be the index of a customer's type of problem, $\mu \in \{l, h\}$ the index of the quality of treatment provided, and $\kappa \in \{l, h\}$ the index of the quality of treatment

charged for. Then the utility of a seller of type (a, b) is assumed to be given by

$$U_{a,b}(p^l, p^h, \theta, \mu) = p^\kappa - c^\mu - aI_{\theta > \mu} - bI_{\theta < \mu}, \quad (3)$$

where $a \geq 0$ is the disutility from undertreating and $b \geq 0$ is the disutility from overtreating a customer. I denotes an indicator variable which takes the value of one if the condition in the subscript is met and the value of zero otherwise.¹⁰

This utility function has straightforward behavioral implications. First, consider condition N. If $c^h - c^l$, i.e. the additional profit a seller receives from providing t^l when t^h is needed, is small compared to the expert's disutility a from undertreating a customer then the expert will provide the appropriate treatment. Overtreatment is never optimal for an expert in N as it is dominated by overcharging (in comparison to overcharging, overtreatment decreases the expert's payoff while leaving the consumer's payoff unaffected). By contrast, in condition V the mark-up difference, i.e. $(p^h - c^h) - (p^l - c^l)$, becomes important for experts' provision behavior. Figure 5 shows the areas of undertreatment and overtreatment when verifiability applies. Important for our analysis is that a taste for efficiency predicts appropriate treatment in a corridor along the equal mark-up line. As is easily verified, this property is an implication of a taste for efficiency *per se*; that is, it does not depend on the specific functional form imposed above (the exact shape of the corridor does, of course).

3.2 Inequality Averse Preferences

Next we analyze the consequences of inequality aversion for the provision behavior on credence goods markets. To illustrate our main points, we assume (like Fehr and Schmidt 1999) that the utility of a seller of type (α, β) does not only depend on her own monetary payoff, π_s , but also on the payoff of the consumer, π_c :

$$U_{\alpha,\beta}(\pi_s, \pi_c) = \pi_s - \alpha(\max\{\pi_c - \pi_s, 0\}) - \beta(\max\{\pi_s - \pi_c, 0\}). \quad (4)$$

We also assume that $\alpha \geq \beta \geq 0$, i.e. a person suffers from inequality and it suffers more from disadvantageous inequality than from inequality in the person's favor; and $\beta < 1$, i.e. the seller refrains from wasting money to

¹⁰Again, we use the convention that $l \leq h$, but not vice versa.




Figure 5: Provision Behavior of an Efficiency Loving (EL) Expert under Verifiability

reduce advantageous inequality because the direct effect on π_s is stronger than the reduced disutility due to a more equal outcome. This is exactly the two-person version of Fehr and Schmidt (1999).

Our main focus is on the case where verifiability holds. For simplicity of presentation, we concentrate on the case where undertreatment (i.e. providing t^l when t^h is needed) implies an outcome where the monetary payoff of the expert exceeds that of the consumer. This restriction seems quite natural and it translates in our framework to a lower bound on p^l , namely to $p^l \geq c^l/2$.¹¹ Note that this condition is satisfied for the parameter values and range restrictions implemented in the experiment reported in the previous section (where $p^l \geq 1$ and $c^l = 2$).

To characterize the provision behavior of an inequality averse expert it is useful to subdivide the space of price vectors illustrated in Figure 6 in four quadrants depending on the sign of the difference in monetary payoffs of the two trading partners in case of appropriate treatment. In Area A the expert's monetary payoff exceeds that of the customer in both cases, when the customer needs t^h and appropriately receives t^h and when he needs t^l and appropriately receives t^l . Thus, this area is defined by $v - p^h < p^h - c^h$ and $v - p^l < p^l - c^l$. In Area B the expert is better off if she appropriately provides high quality, but worse off if she appropriately provides low quality. Hence, this area is defined by $v - p^h < p^h - c^h$ and $v - p^l > p^l - c^l$. In Area C the expert is worse off than the customer in both cases. Thus, this area is defined by $v - p^h > p^h - c^h$ and $v - p^l > p^l - c^l$. Finally, in Area D the expert is worse off if she appropriately provides high quality but better off if she appropriately provides low quality. Thus, this area is defined by $v - p^h > p^h - c^h$ and $v - p^l < p^l - c^l$.

Within each area it is straightforward to solve for the provision behavior of the expert depending on what the consumer needs (details are given in the *supplementary material to this article*). Figure 6 shows the combinations of prices that induce an expert to provide the appropriate treatment, to always provide high quality (overtreatment if the customer needs t^l), to always provide low quality (undertreatment if the customer needs t^h) and to always provide exactly the wrong quality (overtreatment if the customer needs t^l AND undertreatment if the customer needs t^h). The figure reveals that inequality aversion can lead to a positive or a negative deviation from

¹¹As is easily verified, allowing for $p^l < c^l/2$ does not change any of our results - it only complicates the analysis by making it necessary to consider more (sub-) cases.

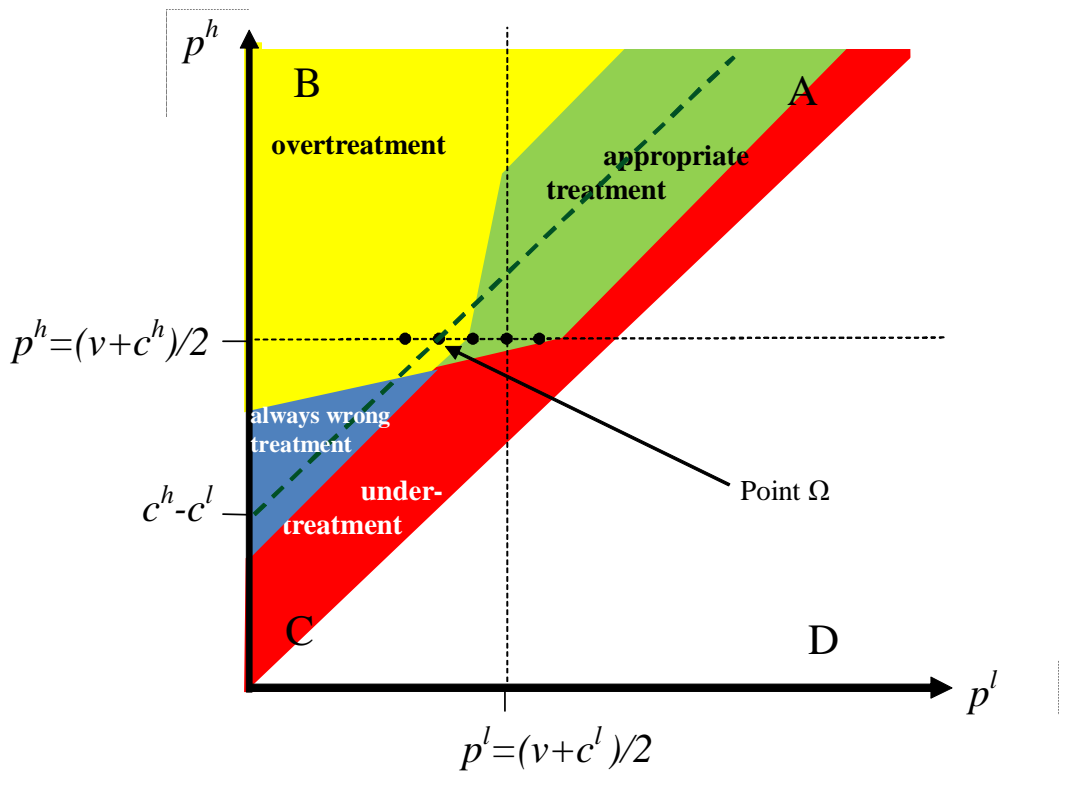


Figure 6: Provision Behavior of an Inequality Averse (IA) Expert under Verifiability

standard theory’s prediction, depending on the price vector under which a transaction takes place: For some price vectors inequality aversion predicts appropriate treatment (to reduce advantageous inequality) when own-money maximization calls for over- or undertreatment (especially in Area A, but also in parts of areas B and D). For other price vectors inequality aversion calls for over- and/or undertreatment (to reduce disadvantageous inequality) when own-money maximization is consistent with providing the appropriate treatment (especially in Area C, but also in parts of areas B and D). Also notice that there is a subarea (of Area C; see Footnote 12 on this) where the expert always provides the wrong treatment, a behavior that is (generically) inconsistent with standard preferences. Here it is important to note that the provision behavior of inequality averse experts is qualitatively as depicted in the figure independently of the magnitude of the parameters α and β .¹² We will argue below that the qualitative features of experts’ provision behavior and -most importantly- our experimental design in the next section do not depend on the linearity assumption either. That is, these features are implied by inequality aversion *per se* and not by the specific functional form imposed above.

3.3 A Test for Discriminating Between Different Types of Distributional Preferences

In this subsection we propose a test for the identification of distributional preferences that relies on a small set of plausible assumptions. The assumptions are that subjects (in the role of experts in the experiment) have preferences that can be represented by a utility or motivation function of the form $U(\pi_s, \pi_c)$ satisfying the following three conditions:

- $\partial U / \partial \pi_s > 0$

¹²In the (linear) Fehr and Schmidt model the picture always looks qualitatively like Figure 6 for all admissible values of α and β , the only exception being that for high values of α combined with low values of β –the exact condition being $\alpha > \beta[v - (c^h - c^l)] / [(c^h - c^l) - 2v\beta]$ – the *always wrong treatment* region enters Area B (implying that the *always appropriate treatment* region disappears from Area C). Note that the *always wrong* region is necessarily to the lower left of the *always appropriate treatment* region and that those two regions necessarily intersect in exactly one point. Also note that the intersection necessarily occurs to the right of the equal mark-up line. More details are available in the *supplementary material to this article*.

- $sign(\partial U/\partial\pi_c)$ depends (only) on whether $\pi_s > \pi_c$, or $\pi_s < \pi_c$
- $\partial U/\partial\pi_s > \partial U/\partial\pi_c$

The first of those assumptions is innocuous, it requires only that – holding the material payoff of the customer constant – the seller’s utility increases in own material payoff. The second assumption is both permissive and restrictive, depending on the perspective. It is permissive because it allows for all variants of distributional preferences that have been discussed in the economics literature, including taste for efficiency, inequality aversion, competitive (or spiteful, or status-seeking) preferences, maximin, quasi-maximin etc. The second assumption is also restrictive because it implies that preferences only depend on outcomes, not on the way these are achieved. Among others, this excludes reciprocity or other intentional motives to play a role (as is the case in all purely outcome-based models). The third assumption is fairly plausible for allocations with $\pi_s < \pi_c$ (it seems sensible to assume that the material payoff of the customer does not have more weight in the expert’s utility function than her own material payoff when the customer is already ahead) but might be regarded as restrictive for allocations with $\pi_s > \pi_c$. This assumption’s main purpose is to ease the exposition, though, and it can be relaxed without changing results qualitatively.¹³ We call an expert :

- **SE - selfish** iff a) $\partial U/\partial\pi_c = 0$ for $\pi_s > \pi_c$; and b) $\partial U/\partial\pi_c = 0$ for $\pi_s < \pi_c$;
- **IA - inequality averse** iff a) $\partial U/\partial\pi_c \geq 0$ for $\pi_s > \pi_c$; b) $\partial U/\partial\pi_c \leq 0$ for $\pi_s < \pi_c$; and c) at least one of the two weak inequalities is strict;
- **EL - efficiency loving** iff a) $\partial U/\partial\pi_c \geq 0$ for $\pi_s > \pi_c$; b) $\partial U/\partial\pi_c \geq 0$ for $\pi_s < \pi_c$; and c) at least one of the two weak inequalities is strict;
- **CO - competitive** iff a) $\partial U/\partial\pi_c \leq 0$ for $\pi_s > \pi_c$; b) $\partial U/\partial\pi_c \leq 0$ for $\pi_s < \pi_c$; and c) at least one of the two weak inequalities is strict;
- **IL - inequality loving** iff a) $\partial U/\partial\pi_c \leq 0$ for $\pi_s > \pi_c$; b) $\partial U/\partial\pi_c \geq 0$ for $\pi_s < \pi_c$; and c) at least one of the two weak inequalities is strict.

¹³Technically the purpose of the assumption is to get a unique ”switching point” in the test proposed below.

A selfish (SE) seller is a homo oeconomicus according to standard theory, she simply maximizes her own material payoff. An inequality averse (IA) expert wants to see the payoff of her customer increased if she is better off than the customer; but she wants to see the customer's payoff decreased if the opposite is the case. An efficiency loving (EL) expert is willing to give up own monetary payoff to increase the material payoff of her trading partner if the 'price of giving' is not too high. A competitive (CO) expert is willing to give up own material payoff to decrease the payoff of her trading partner if the 'price of taking' is not too high.¹⁴ An inequality loving (IL) expert is willing to sacrifice own material payoff to increase the difference between the payoffs of the two trading partners.¹⁵ Figure 7 displays typical indifference curves of the different types of distributional preferences in the own-payoff/other-payoff space.

How can we discriminate between the different types of distributional preferences? In the space of possible price vectors there is exactly one that allows for a neat discrimination. Looking at Figure 6 it is the price vector at the intersection of the upward sloping dashed line and the horizontal dashed line. For future reference we name (the price vector lying at) this intersection 'Point Ω '.

Consider Point Ω and assume that the expert is inequality averse. For this case Figure 6 (based on the assumption that experts have Fehr/Schmidt preferences) predicts overtreatment. This is not only true for a certain para-

¹⁴A competitive (or spiteful) seller has an incentive to hurt her customer if this is not too costly for her. Thus, a representation of her provision behavior in the (p^l, p^h) space would yield a picture similar to the one for the efficiency loving seller (Figure 5) except that the corridor along the equal mark-up does not lead to appropriate treatment but to *always wrong* provision (that is, a customer with the minor problem gets the major treatment, while a customer with the major problem gets the minor treatment). In other words, the green area in Figure 5 would become blue if experts have competitive preferences, implying that no ("green") price vectors exist that induce experts of this type to provide the appropriate treatment.

¹⁵Inequality loving (or equality averse) types are included for completeness only; we do not expect to find many of them (even though Fershtman, Gneezy and List 2009, show that inequality loving can prevail in competitive settings). A representation of the provision behavior of an IL type in the (p^l, p^h) space would yield a picture similar to the one for the inequality averse seller (Figure 6) except that (i) the green area does not lead to appropriate treatment but to *always wrong* provision, (ii) the blue area does not lead to *always wrong* provision but to appropriate treatment, and (iii) the intersection point between the green and the blue area does not occur to the right but rather to the left of the equal mark-up line.

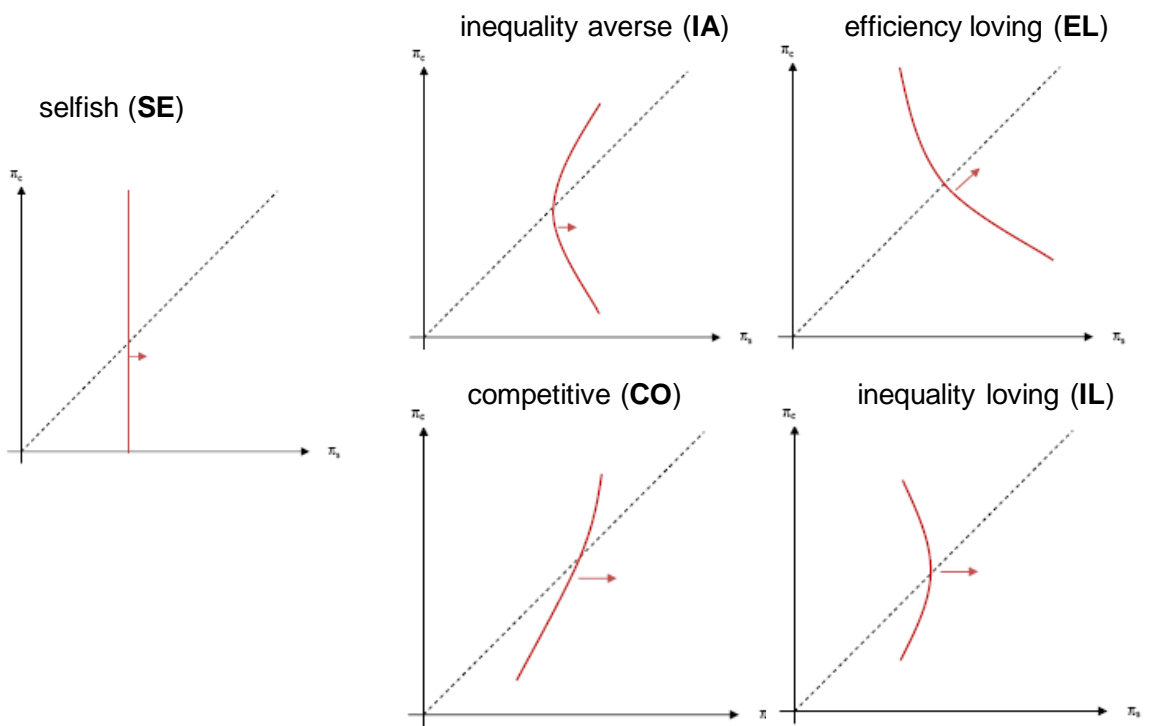


Figure 7: Indifference Curves of SE, IA, EL, CO and IL Types in the Own Payoffs (Horizontal Axis) and Other-Player's Payoffs (Vertical Axis) Space

meter range in the Fehr and Schmidt model, it is not even only true for the specific (linear) functional form, it is rather implied by the spirit of inequality aversion *per se*. To see this point, first notice that along the upward sloping dashed line (the equal mark-up line) we have $p^h - p^l = c^h - c^l$, while along the horizontal dashed line we have $\pi_c(p^l, p^h, \theta = h, \mu = h, \kappa = h) = \pi_s(p^l, p^h, \mu = h, \kappa = h) \iff p^h = (v + c^h)/2$. In words, along the upward sloping dashed line the expert gets exactly the same material payoff independently of whether she provides t^l or t^h , while along the horizontal dashed line the expert and the customer get exactly the same material payoff if the expert provides t^h . Thus, at the intersection of the two lines in Point Ω we have $p^l = (v + c^l)/2 - (c^h - c^l)/2$ implying that this point is necessarily to the left of the vertical dashed line where we have $\pi_c(p^l, p^h, \theta = l, \mu = l, \kappa = l) = \pi_s(p^l, p^h, \mu = l, \kappa = l) \iff p^l = (v + c^l)/2$.

Now suppose we (as the experimentalists) impose the price vector in Point Ω and look at an expert's provision behavior. First assume the customer has the minor problem. If the expert provides the minor treatment, she induces a payoff allocation $(\pi_s, \pi_c) = (x, y)$ with $x < y$, where the strict inequality follows from the fact that Point Ω is (strictly) to the left of the vertical dashed line. If the expert provides the major treatment instead, she induces a payoff allocation $(\pi_s, \pi_c) = (x_1, x_2)$ with $x_1 = x$ (follows from the fact that Point Ω is on the upward sloping dashed line) and $x_2 = x$ (follows from the fact that Point Ω is on the horizontal dashed line). Thus, inequality aversion dictates deciding for the second allocation (that is, providing the major treatment to a customer who has the minor problem) because the own material payoff is the same in both allocations while disadvantageous inequality is present in the first but absent in the second allocation. Now assume that the customer has the major problem. If the expert provides the major treatment, she induces a payoff allocation $(\pi_s, \pi_c) = (x_3, x_4)$ with $x_3 = x_1 = x$ and $x_4 = x_2 = x$, where the equalities follow from the fact that the material payoffs (of both parties) for providing the major treatment are independent of the type of problem. If the expert provides the minor treatment instead, she induces a payoff allocation $(\pi_s, \pi_c) = (x_5, z)$ with $x_5 = x_1 = x$ (follows from the fact that Point Ω is on the upward sloping dashed line) and $x_5 = x > z$ (from $p^l \geq c^l/2$). Thus, inequality aversion dictates deciding for the first allocation (that is, providing the major treatment to a customer who has the major problem) because the own material payoff is the same in both allocations while advantageous inequality is present in the second allocation but absent in the first. In sum: Inequality aversion *per se* (and not only a specific version

of it) calls for overtreatment in Point Ω .

Performing a similar analysis for the other types of distributional preference and noting that $y > x > z$ yields the following prediction:

Prediction 2a (Impartial Distribution Preferences) *Consider the price vector Ω with $p^h = (v + c^h)/2$ and $p^h - c^h = p^l - c^l$ (and note that the two equations together imply $p^l < (v + c^l)/2$). Under this price vector:*

- a) *appropriate treatment is consistent with selfish preferences and **efficiency loving** preferences but inconsistent with IA, CO and IL;*
- b) *overtreatment is consistent with selfish preferences and **inequality aversion** but inconsistent with EL, CO and IL;*
- c) *undertreatment is consistent with selfish preferences and **inequality loving** preferences but inconsistent with IA, CO and EL;*
- d) *always wrong treatment is consistent with selfish preferences and **competitive** preferences but inconsistent with IA, EL and IL.*

Testing the provision behavior under the price vector described in Prediction 2a is like eliciting impartial distributional preferences, because at Point Ω a seller compares two allocations that yield the same material payoff for her, but different payoffs for the customer. Thus, deciding for the "fair" allocation (whatever is considered fair) does not involve any costs. In the test suggested in the next section we start with the price vector in Point Ω and then we change p^l slightly, keeping p^h constant. In Figure 6 this means that we move along the horizontal dashed line. In terms of payoff allocations, moving to the right (left, respectively) of Point Ω means that we increase (decrease) x at the cost (for the benefit) of y and x_5 at the cost (for the benefit) of z , while keeping $x_1 = x_2 = x_3 = x_4$ constant. Given our three assumptions on the utility or motivational function $U(\pi_s, \pi_c)$, what are the implications of such a change for the provision behavior of sellers with different types of distributional preferences? First, it implies a kind of *monotonicity*:

Prediction 2b (Monotonicity) *Consider two price vectors, Ω and Ψ . Suppose that both have the same p^h given by $p_\Omega^h = p_\Psi^h = (v + c^h)/2$ but different p^l ($p_\Omega^l \neq p_\Psi^l$) with $p_\Omega^l = (v + c^h)/2 - (c^h - c^l)$ and $p_\Psi^l < (v + c^l)/2$.*

- *If $p_\Omega^l < p_\Psi^l$ then - keeping the consumer's type of problem constant - an expert who provides t^l under Ω must provide t^l under Ψ .*
- *If $p_\Omega^l > p_\Psi^l$ then - keeping the consumer's type of problem constant - an expert who provides t^h under Ω must provide t^h under Ψ .*

Proof. First note that providing t^h yields the equal material payoff allocation $\pi_s = \pi_c = (v - c^h)/2$ independently of the consumer's type of problem and independently of whether Ω or Ψ is the relevant contract. By contrast, the payoff allocation from providing t^l depends on both, the type of problem and the type of contract. Suppose first the *consumer needs t^l* . Then providing t^l under Ω yields $\pi_s = (v - c^h)/2$ and $\pi_c = (v - c^h)/2 + (c^h - c^l)$, while providing t^l under Ψ yields $\pi_s = (v - c^h)/2 + \varepsilon$ and $\pi_c = (v - c^h)/2 + (c^h - c^l) - \varepsilon$, where $\varepsilon > 0$ for $p_\Omega^l < p_\Psi^l$ and $\varepsilon < 0$ for $p_\Omega^l > p_\Psi^l$. Now suppose the *consumer needs t^h* . Then providing t^l under Ω yields $\pi_s = (v - c^h)/2$ and $\pi_c = (v - c^h)/2 + (c^h - c^l) - v$, while providing t^l under Ψ yields $\pi_s = (v - c^h)/2 + \varepsilon$ and $\pi_c = (v - c^h)/2 + (c^h - c^l) - v - \varepsilon$, where $\varepsilon > 0$ for $p_\Omega^l < p_\Psi^l$ and $\varepsilon < 0$ for $p_\Omega^l > p_\Psi^l$. It remains to be shown that $U(\varphi + \varepsilon, \chi - \varepsilon)$ is increasing in ε . This follows from $\partial U/\partial \pi_s > \partial U/\partial \pi_c$ for all (φ, χ) ■

Prediction 2a and Prediction 2b together imply the following result:

Prediction 2c (Partial Distribution Preferences) *Consider the price vectors Ω and Ψ as defined in Prediction 2b. Then observing*

- a) appropriate treatment under Ω and Ψ is only consistent with **efficiency loving** preferences (but inconsistent with SE, IA, CO and IL);*
- b) overtreatment under Ω and overtreatment, appropriate treatment or always wrong treatment under Ψ with $p_\Omega^l < p_\Psi^l$ is only consistent with **inequality aversion** (but inconsistent with SE, EL, CO and IL);*
- c) undertreatment under Ω and undertreatment, appropriate treatment or always wrong treatment under Ψ with $p_\Omega^l > p_\Psi^l$ is only consistent with **inequality loving** preferences (but inconsistent with SE, IA, CO and EL);*
- d) always wrong treatment under Ω and always wrong under Ψ is only consistent with **competitive** preferences (but inconsistent with SE, IA, EL and IL).*

To understand Prediction 2c (and the term 'partial' distribution preferences) consider an inequity averse seller. From the arguments above we know that such an expert has to overtreat a customer under price vector Ω . Increasing p^l slightly while keeping p^h constant creates a tension between a higher own monetary payoff and more inequality and vice versa for an inequality averse expert. By deciding for overtreatment or switching to appropriate treatment (or always wrong treatment) she reveals a positive willingness to pay for reducing inequality. The argument for sellers with other kinds of distributional preferences is similar.

4 Testing for the Presence/Absence of Non-Standard Preferences

To test for and classify the distributional preferences of sellers, we conducted a new series of credence goods experiments. In these new experiments, the timing of the game is exactly the same as in the game described in Section 2 except for the first stage: Instead of letting sellers post their prices themselves (which might give rise to endogeneity concerns), we let the (exogenous) experimental software choose a price vector from the set $\{(3,8), (4,8), (5,8), (6,8), (7,8)\}$ with equal probability in each round. This set of vectors has two characteristics: First and foremost, it includes the equal mark-up vector Ω characterized in Prediction 2a; and starting from this equal mark-up vector it varies p^l as described in Prediction 2b and Prediction 2c. Second, this set of price-vectors includes the four most frequently chosen price vectors in condition V, where prices were endogenous (see Table 2). We call the experimental condition with this (exogenously given) set of price vectors V-Fix1. In order to check whether the inclusion of the price vector (3,8) - which had been chosen in condition V in only 2 out of 704 cases - would have any impact on behavior we also ran another experimental condition where the exogenously determined price vector was chosen with equal probability only from the top four price vectors (4,8), (5,8), (6,8), and (7,8). We call this condition V-Fix2. We ran four sessions with 16 subjects each both in V-Fix1 and V-Fix2. Hence, a total of 128 subjects participated in these new sessions (with no subject having participated in the sessions reported in Section 2). Sessions lasted less than 1.5 hours and average earnings were about 15 Euro.

4.1 Aggregate Behavior

As before we start with some details on aggregate behavior. Table 5 reports the under- and overtreatment rates for the five price vectors. Since we do not find any significant differences in undertreatment and overtreatment rates between V-Fix1 and V-Fix2 we provide pooled data from both treatments.

Observation 3 (Aggregate Behavior in Condition V-Fix) *First, the undertreatment and the overtreatment rate are rather high under the equal mark-up vector (4,8). Second, there is a considerable amount of appropriate treatment under each price vector. Third, monetary incentives work roughly*

in the direction standard theory predicts: an increase in p^l holding p^h constant increases the undertreatment rate and decreases the overtreatment rate. Fourth, there is a non-monotonicity at (6,8): at this price vector both the undertreatment and the overtreatment rate are (slightly) too low to yield monotonicity.

Table 5: Under- and Overtreatment Rates in Condition V-Fix

(p^l, p^h)	Undertreatment Rate	Overtreatment Rate
(3,8)	3.57%	93.48%
(4,8)	22.11%	41.18%
(5,8)	64.10%	5.26%
(6,8)	63.23%	3.85%
(7,8)	66.07%	4.55%

How can the aggregate behavior be explained? Examining the experimental results on the individual level helps to answer this question.

4.2 Classifying Individual Behavior

Observation 4 (Individual Behavior in Conditions V-Fix) *Less than a fourth of the experimental sellers act according to standard theory's prediction: they provide appropriate treatment if and only if they are held indifferent in own-money terms. About a fourth of the seller population displays behavior that is consistent with a strong taste for efficiency: they provide appropriate treatment even if own-money maximization calls for over- or undertreatment. About another fifth of sellers shows behavior that is consistent with strong inequality aversion: they overtreat customers even if this reduces their own monetary payoff. Adding up strong and weak forms of distributional preferences indicates that about half of the sellers display behavior that is consistent with a taste for efficiency, while little more than a fourth of the sellers display behavior consistent with (strong or weak) inequality aversion.*

In the data analysis of individual sellers' preference types we first looked at violations of monotonicity according to Prediction 2b. It turns out that 45 out of 64 sellers (70%) behave according to the prediction over all 16 periods of the experiment. Taking into account that some learning may go on in early periods, we decided to look in the following at the final 12 periods only

(i.e., periods 5 to 16) and found that the behavior of 56 out of 64 sellers (88%) respects the monotonicity condition as defined in Prediction 2b. This high degree of consistent behavior is encouraging because it suggests that stable (non-standard) preferences, rather than noise or any kind of confusion of subjects, drives our findings. Of the 56 sellers whose behavior is consistent with Prediction 2b, we had to exclude 3 from further analysis due to lack of data caused by customers' opting out.¹⁶ Table 6 is therefore based on 53 sellers.

Table 6: Classification of Individual Behavior in V-Fix

Type of Behavior	strong	weak	total
EL (efficiency loving)	14 (26.4%)	12 (22.6%)	26 (49.0%)
IA (inequality averse)	10 (18.9%)	4 (7.5%)	14 (26.4%)
CO (competitive)	0 (0%)	3 (5.7%)	3 (5.7%)
IL (inequality loving)	0 (0%)	2 (3.8%)	2 (3.8%)
SE (selfish)	8 (15.1%)	21 (39.6%)	29 (54.7%)

To read Table 6 properly, note that sellers who are classified as either weak EL, weak IA, weak CO, or weak IL types are also classified as weak SE types. This has to be the case because weak EL, IA, CO and IL types behave exactly as the strong version of the respective type as impartial arbitrators (that is, when there is no trade off between own material payoff and fairness standard as at Point Ω in Figure 6), but they behave exactly like (strong) SE types when their own material payoff is at stake.¹⁷ Thus, for relative frequencies (given in parentheses in Table 6) to add up to 100%, one has to

¹⁶Recall that we only changed the first stage of the game in the V-Fix conditions. The second stage, where customers could decide whether or not to interact with a seller – at exogenously given prices – was also present in the V-Fix conditions. The 3 sellers that we had to exclude from the further analysis had too few interactions to assign them to one particular type. Our criteria for inclusion/exclusion were as follows: We included all experts who have treated under price vector Ω at least one customer needing t^l AND at least one customer needing t^h . 50 of the 56 sellers were included under this rule. From the remaining 6 sellers, we included those where the data contained a clear indication that the expert is of one of the major distributional types (and not of one of the others). Since only 3 sellers were included under this latter rule and since only 6 seller were at disposition a change in the criteria of inclusion/exclusion would not have changed our results qualitatively.

¹⁷Formally, the reason is that the weak SE type is the limit of all kinds of distributional types "when the weight on the distributional part of the utility function goes to zero". Note, however, that the limiting behavior is different for the four non-SE types!

add either the strong non-SE types and the total number of SE types or the total number of non-SE types and the number of strong SE types.

An important point is that the behavior of only a minority of individuals is consistent with standard theory's assumption that sellers always follow their monetary incentives and in case of indifference they act in the interest of customers. Note that less than a fourth of experimental sellers (those in the category "weak EL") act in this manner. This provides an explanation for both, why equal mark-up vectors do not work as predicted by theory, and why these price vectors were not chosen in the endogenous pricing conditions.

Another point that can be inferred from Table 6 is that individuals in the lab are heterogeneous.¹⁸ Some sellers care for efficiency, some for 'equality of payoffs', some do not care for the well-being of others (or for efficiency) at all. An important implication is that while studying aggregate data can give first evidence and provide ideas for further elaboration, analyzing individual behavior is important to understand the trade-offs that influence behavior of subjects on experimental credence goods markets.

5 Conclusions

This paper has started from the observation that - contrary to standard theory's predictions - verifiability has no influence on key variables on an experimental credence goods market: While (standard) theory predicts that verifiability brings full efficiency whereas the absence of it leads to a market break down, we have found that the frequency of interaction, the undertreatment rate, the overtreatment rate and overall efficiency are not significantly different in credence goods markets with and without verifiability. The overall performance in both conditions has been better than the standard prediction for markets where verifiability is violated, but worse than the standard prediction for markets where verifiability holds. Furthermore, an analysis of individual behavior has indicated that non-standard distributional preferences shape the provision behavior of experts.

Motivated by these findings, we have studied the implications of different types of distributional preferences for sellers' provision and charging behavior in credence goods markets under verifiability. We have found that the impact

¹⁸Of course, the same holds true in many other games, such as public goods games (Fischbacher, Gaechter and Fehr 2001; Fischbacher and Gächter 2009) or gift-exchange games (Fehr et al. 1997).

on market efficiency can be either positive or negative, depending on the type of distributional preferences and the price vector under which trade takes place. While given prices may induce some preference types to behave appropriately, the same prices may induce other types to choose always the most inefficient strategy. Based on this observation we have identified a single price vector that allows a clean discrimination between different preference types using only a minimal set of assumptions. An important feature of the identified price vector is that the discrimination does not depend on any functional form or parameterization of the underlying utility or motivational function. The price vector rather directly tests the key characteristics of different variants of distributional preferences that have been discussed in the literature.

Based on our theoretical analysis we have then designed an experiment that allows for a classification of various distributional preference-types. The key findings of this experiment (with exogenously determined prices) are the following: (i) Only a minority (of about 20%) of individuals behave according to standard theory's prediction: they provide appropriate treatment if and only if they are held indifferent in own-money terms. (ii) About half of experimental sellers care for efficiency. (iii) Roughly one quarter of sellers cares for 'equality of payoffs' of sellers and buyers. (iv) A minority of experts (of about 5%) have competitive preferences, meaning that they have a propensity to harm their customers.

Taken together, these findings provide an explanation for both, why verifiability does a poor job in securing efficiency on credence goods markets and why markets without verifiability perform considerably better than predicted. The key insight to see this point is that one and the same heterogeneity in distributional preferences can lead to a positive or a negative deviation from standard theory's prediction depending on institutional design and - within a given design - depending on the price vector under which the transaction takes place. The positive deviation may arise from efficiency loving or inequality averse experts. Efficiency loving experts are willing to provide the appropriate treatment if the personal cost of doing so is not too high, while inequality averse experts care for their customers (at a reasonable cost) in the domain of advantageous inequality. The negative deviation potentially comes from inequality averse or competitive experts. Inequality averse experts have a propensity to harm their customers in the domain of disadvantageous inequality, while competitive experts are willing to hurt customers if the 'price for hurting' is not too high.

We will argue in the following that the institutional design N is robust against the negative impact of distributional preferences, but not against the positive one, while under V the opposite is true. It is important to note that our arguments do not depend on any assumptions on the distribution of different social-preference-types in the population (i.e., on which fraction of subjects exhibits inequality aversion, which fraction a taste for efficiency, etc.) and on how specific social preferences are modelled. What is important for our arguments, though, is that there is heterogeneity in preference types.

Consider *condition N* first. Under this condition the negative consequences of distributional preferences cannot manifest themselves in the market outcome since the standard prediction (undertreatment and overcharging under each price vector) is already a worst case scenario that leaves no room for deterioration. In other words, experts with negative attitudes towards customers behave exactly like own-money maximizing experts in condition N. However, the positive side of distributional preferences immediately manifests itself in a better market outcome than predicted under standard preferences in the sense that experts provide appropriate treatment despite strong monetary incentives for undertreatment.

Under *condition V*, by contrast, we get almost the opposite result (where the ‘almost’ does not apply along standard theory’s equilibrium path where experts are predicted to post equal mark-up prices). Under equal mark-up price vectors the positive effects of distributional preferences cannot manifest themselves in the market outcome as the standard prediction (appropriate treatment independent of the level of mark-ups) is already a best case scenario. However, the negative effects of distributional preferences easily manifest themselves in the market outcome because hurting the customer involves no cost under equal mark-up vectors. This argument does not only explain the poor performance of equal mark up prices in condition V, it also explains why equal mark up prices are very rarely chosen in the first place.

Taking the evidence from the N-condition and the V-condition together we see that the net effect of distributional preferences is positive in condition N, but negative in condition V, leading to both conditions performing equally (well or poor) in the aggregate.¹⁹

¹⁹As we saw in Section 2, more than 90% of the price vectors in the endogenous price choice condition V were of the undertreatment type. The impact of distributional preferences under undertreatment vectors in condition V is quite similar to the impact of distributional preferences (under any price vector) in condition N – the only difference being that the “price for giving” depends on the price difference between qualities in V

While the parameter-free classification of subjects according to their distributional preferences is an interesting and important contribution of the present paper in itself, important conclusions for credence goods-markets and, more generally, for markets with asymmetric information can be drawn from our analysis. An immediate implication is that institutional design based on the standard assumption of own-money-maximizing subjects yields bad incentives for a majority of agents. Another implication is that there are agents that behave appropriately independently of the institutional design. Taken together these two observations have two important consequences, one for institutional design, the other for agent selection:

Designing the Right Institutions: What is needed are not perfect institutions for one type of agent, but rather institutions that are robust against the coexistence of different types of agents. Our results clearly show that verifiability is not such an institution (nor is a market where verifiability does not apply). By contrast, as Dulleck et al. (2009) have shown, 'liability' is a quite robust institution in markets for credence goods.²⁰ 'Liability' requires verifiability of 'outcomes', while 'verifiability' requires only verifiability of 'inputs'. Thus, securing verifiability of outcomes, where possible, might solve credence goods problems more effectively.

Selecting the Right Agents: Designing robust institutions might be difficult, though. Imposing liability, for instance, generates other problems or may be impossible to achieve.²¹ As a consequence, selecting the "right"

while it is fixed in N (this difference is a minor one, however, since observed price differences are typically small): The standard prediction is already a worst case scenario that leaves no room for deterioration. So the negative part of distributional preferences cannot have an impact. The positive part, on the other hand, easily manifests itself in the market outcome. This might explain why the overall performance in the endogenous price choice experiments is quite similar in both conditions.

²⁰Under liability the negative effects of distributional preferences cannot lead to deviations from the standard prediction since undertreatment is ruled out by design and since overcharging is already the standard prediction. Hence, only overtreatment remains as a possible harm to efficiency, but overtreatment (in comparison to overcharging) only reduces the expert's material payoff without affecting that of the customer and is therefore unattractive for experts independently of their distributional type (provided the expert's utility or motivational function $U(\pi_s, \pi_c)$ satisfies $\partial U / \partial \pi_s > 0$, which is quite plausible, as argued earlier). The positive effects of distributional preferences cannot have a deep impact either, since the benchmark prediction is already full efficiency – so the only positive impact can be a distributional one leading to overcharging rates below the predicted 100% (and this is indeed what Dulleck et al. 2009 find in their experiments).

²¹On the one hand, liability requires a form of verifiability of the outcome. Especially

agents for jobs involving experts' services becomes particularly important. Instead of choosing the best trained doctor, mechanic or computer specialists, customers or their representatives should also worry about the attitudes of these experts towards their customers. Selecting the right agents may also help to solve problems created by cost uncertainty over inputs: With cost uncertainty standard theory would predict that verifiability cannot solve the problems on credence goods markets - a problem ignored in the formal literature on credence goods thus far. Our results suggest that verifiability can solve this problem if the "right" agents are selected: Efficiency loving experts provide appropriate treatment in a corridor along the equal mark-up line; that is, even if monetary incentives are not perfectly in line. Hence, the crucial task of potential employers or buyers is to identify the experts' distributional, or more generally speaking, social preferences. While a candidate's track record in his CV (through social activities) may provide first clues, repeated interaction will most likely provide additional evidence on an expert's distributional preference-type.

Condensed to a single sentence our results show that good performance of credence goods markets requires either robust institutions (combined with arbitrary experts) or "good" experts (combined with arbitrary institutions).

References

- [1] Alger, I. and Salanié, F. (2006), A Theory of Fraud and Overtreatment in Experts Markets, *Journal of Economics and Management Strategy* **15**, 853-881.
- [2] Bolton, G. E. and Ockenfels, A. (2000), ERC – A Theory of Equity, Reciprocity and Competition. *American Economic Review* **90**: 166-193.

in the medical realm treatment success is often impossible or very costly to measure for a court, while still being observed by the consumer (how can one prove the presence/absence of pain, for instance?). On the other hand, even in cases where the outcome is verifiable (for instance, in the repair business) strict liability might pose problems. For instance, an insufficiently repaired car may work for some time before it breaks down. To mitigate the undertreatment problem in such a situation the liability needs to cover a longer period. But during this longer period the car may stop working for reasons unrelated to the expert's behavior. Also, an extended liability period may induce fraudulent behavior on the side of the customer (as she may not put in the required maintenance effort) in situations where the eventual performance of the product depends not only on the expert's, but also on the consumer's behavior. Taylor (1995) discusses these double-sided incentive problems.

- [3] Charness, G. and Rabin, M. (2002), Understanding Social Preferences with Simple Tests, *Quarterly Journal of Economics*, **117**: 817-869.
- [4] Darby, M. and Karni, E. (1973), Free Competition and the Optimal Amount of Fraud, *Journal of Law and Economics* **16**, 67-88.
- [5] Dulleck, U. and Kerschbamer, R. (2006), On Doctors, Mechanics, and Computer Specialists: The economics of credence goods. *Journal of Economic Literature* 44: 5-42.
- [6] Dulleck, U. and Kerschbamer, R. (2008), Price Discrimination in a Market for Credence Goods, Mimeo, Department of Economics, University of Innsbruck.
- [7] Dulleck, U. and Kerschbamer, R. (2009), Experts vs. Discounters: Competition and Market Unraveling and Consumers Do Not Know What They Need, *International Journal of Industrial Organization* **27**, 15-23.
- [8] Dulleck, U., Kerschbamer, R. and Sutter, M. (2009), The Economics of Credence Goods: On the Role of Liability, Verifiability, Reputation and Competition, forthcoming in *American Economic Review*.
- [9] Emons, W. (1997), Credence Goods and Fraudulent Experts, *Rand Journal of Economics* **28**, 107-119.
- [10] Emons, W. (2001), Credence Goods Monopolists, *International Journal of Industrial Organization* **19**, 375-89.
- [11] Fehr, E., Gächter, S. and Kirchsteiger, G. (1997), Reciprocity as a Contract Enforcement Device. *Econometrica* **65**: 833-860.
- [12] Fehr, E., Kirchsteiger, G. and Riedl, A. (1993), Does Fairness Prevent Market Clearing? An Experimental Investigation, *Quarterly Journal of Economics* **108**, 437-459.
- [13] Fehr, E. and Schmidt, K. (1999), A Theory of Fairness, Competition, and Cooperation. *Quarterly Journal of Economics* **114**: 817-868.
- [14] Fershtman, C., Gneezy, U. and List, J. (2009), Inequality, Selfishness, and Social Norms, Mimeo, Eitan Berglas School of Economics, Tel Aviv University.

- [15] Fischbacher, U. (2007), Z-tree: Zürich Toolbox for Ready-Made Economic Experiments. *Experimental Economics* **10**: 171-178.
- [16] Fischbacher, U. and Gächter, S. (2009), Social Preferences, Beliefs, and the Dynamics of Free Riding in Public Good Experiments. *American Economic Review*, forthcoming.
- [17] Fischbacher, U., Gächter, S. and Fehr, E. (2001), Are People Conditionally Cooperative? Evidence from a Public Goods Experiment. *Economic Letters* **71**: 397-404.
- [18] Greiner, B. (2004), An Online Recruitment System for Economic Experiments. In: Kremer, K. and Macho, V. (eds.): *Forschung und wissenschaftliches Rechnen 2003*. GWDG Bericht 63. Gesellschaft für Wissenschaftliche Datenverarbeitung, Göttingen, 79-93.
- [19] Gruber, J., Kim, J. and Mayzlin, D. (1999), Physician Fees and Procedure Intensity: The Case of Cesarean Delivery, *Journal of Health Economics* **18**, 473-490.
- [20] Gruber, J. and Owings, M. (1996), Physician Financial Incentives and Cesarean Section Delivery. *RAND Journal of Economics* **27**: 99-123.
- [21] Huck, S., Lünser, G. and Tyran, J.-R. (2006), Competition Fosters Trust. UCL Working Paper.
- [22] Huck, S., Lünser, G. and Tyran, J.-R. (2007), Pricing and Trust. UCL Working Paper.
- [23] Hughes, D. and Yule, B. (1992), The Effect of Per-item Fees on the Behavior of General Practitioners, *Journal of Health Economics* **11**, 413-437.
- [24] Pesendorfer, W. and Wolinsky, A. (2003), Second Opinions and Price Competition: Inefficiency in the Market for Expert Advice, *Review of Economic Studies* **70**, 417-437.
- [25] Schneider, H. (2006), A Field Experiment to Measure Agency Problems in Auto Repair, Mimeo, Yale University.
- [26] Taylor, C. (1995), The Economics of Breakdowns, Checkups, and Cures, *Journal of Political Economy* **103**, 53-74.

- [27] Wolinsky, A. (1993), Competition in a Market for Informed Experts' Services, *Rand Journal of Economics* **24**, 380-398.
- [28] Wolinsky, A. (1995), Competition in Markets for Credence Goods, *Journal of Institutional and Theoretical Economics* **151**, 117-131.

Supplementary Material for:
The Impact of Distributional Preferences on
(Experimental) Markets for Expert Services

Rudolf Kerschbamer
University of Innsbruck*

Matthias Sutter
University of Innsbruck & University of Gothenburg[†]

Uwe Dulleck
Queensland University of Technology[‡]

14. December 2009

Abstract

This supplement characterizes the provision behavior of Fehr and Schmidt (1999) inequality-averse experts in the (p^l, p^h) -space (Figure 6 in the main article), derives comparative statics with respect to α and β , and discusses how the linear Fehr&Schmidt model can be extended to draw figures analogous to Figure 6 for (linear versions of) other forms of distributional preferences.

*rudolf.kerschbamer@uibk.ac.at

†matthias.sutter@uibk.ac.at

‡uwe.dulleck@qut.edu.au

1 Introduction

Figure 6 is based on the two person version of Fehr and Schmidt (1999)'s model of inequality aversion. Applied to the present context of a credence goods market, this model assumes that the utility of a seller of type (α, β) does not only depend on her own monetary payoff π_s but also on the payoff of the consumer, π_c :

$$U_{\alpha,\beta}(\pi_s, \pi_c) = \pi_s - \alpha(\max\{\pi_c - \pi_s, 0\}) - \beta(\max\{\pi_s - \pi_c, 0\}), \quad (1)$$

where $\alpha \geq \beta \geq 0$ and $\beta < 1$. To characterize the provision behavior of an inequality averse expert in the (p^l, p^h) -space it is useful to subdivide this space in 4 areas depending on the sign of the difference in monetary payoffs of the two trading partners in case of appropriate treatment (as it is done in the main article). In Area A the expert's monetary payoff exceeds that of the customer in both cases, when the customer needs t^h and appropriately receives t^h and when he needs t^l and appropriately receives t^l . Thus, this area is defined by $v - p^h < p^h - c^h$ and $v - p^l < p^l - c^l$. In Area B the expert is better off if she appropriately provides high quality but worse off if she appropriately provides low quality. Hence, this area is defined by $v - p^h < p^h - c^h$ and $v - p^l > p^l - c^l$. In Area C the expert is worse off than the customer in both cases, when the customer needs t^h and appropriately receives t^h and when he needs t^l and appropriately receives t^l . Thus, this area is defined by $v - p^h > p^h - c^h$ and $v - p^l > p^l - c^l$. Finally, in Area D the expert is worse off if she appropriately provides high quality but better off if she appropriately provides low quality. Thus, this area is defined by $v - p^h > p^h - c^h$ and $v - p^l < p^l - c^l$.

For future reference we define

$$\bar{p}^h = \frac{v + c^h}{2}$$

and

$$\bar{p}^l = \frac{v + c^l}{2}$$

2 The Four Areas and Some Comparative Statics on Provision Behavior

In what follows many paragraphs start with a code of the form (XYZ). This code refers to the condition that is derived in the paragraph. The first letter indicates the area considered - A,B,C or D - and the last two letters indicate that either the no undertreatment (NU) or the no overtreatment (NO) condition is derived. For example, (CNU) is the paragraph deriving the no undertreatment condition for Area C. The critical value of p^h determined by the condition is stated as a function of p^l , for example $p_{CNU}^h(p^l)$.

In **Area A** (blue lines in figures S1 and S2 below) the seller is better off than the customer if she serves the customer appropriately, independently of the type of treatment the customer needs.

(**ANU**) The expert will not undertreat (blue dashed line) iff $p^h - c^h - \beta(2p^h - c^h - v) \geq p^l - c^l - \beta(2p^l - c^l)$ or

$$p^h \geq p_{ANU}^h(p^l) = p^l + \frac{(1 - \beta)(c^h - c^l) - \beta v}{1 - 2\beta}.$$

How does the critical value $p_{ANU}^h(p^l)$ change in α and β ? Since in Area A all decisions lead to advantageous inequality for the seller, $p_{ANU}^h(p^l)$ is independent of α . With respect to β : $\frac{\partial}{\partial \beta}(p_{ANU}^h(p^l)) = -\frac{(v - c^h + c^l)}{(2\beta - 1)^2} < 0$, where the last inequality follows from $v > c^h$ (for $v < c^h$ provision of t^h would be inefficient). Furthermore, for $\beta = 0$, $p_{ANU}^h(p^l)$ coincides with the equal mark-up line ($p^h = p^l + c^h - c^l$). Together these findings imply that $p_{ANU}^h(p^l)$ lies below and to the right of the equal mark-up line for all positive values of α and β . Referring to Figure S1 note that the black arrows originating from the dashed blue line and from Point Y show the comparative statics with respect to β .

In **Figure S1** and **Figure S2** we use the following **convention**: The direction of a **black arrow** shows the effects of an increase in β , while the direction of a **red arrow** shows the effects of an increase in α . The **absence of a black (red) arrow** indicates that the respective $p_{ABC}^h(p^l)$ or the location of a point does not depend on β (α respectively).

(ANO) The expert will not overtreat (blue solid line) iff: $p^l - c^l - \beta(2p^l - c^l - v) \geq p^h - c^h - \beta(2p^l - c^l - v)$ or

$$p^h \leq p_{ANO}^h(p^l) = p^l + \frac{(1 - \beta)(c^h - c^l)}{1 - 2\beta}.$$

Again this condition does not depend on α and the comparative statics with respect to β are $\frac{d}{d\beta}(p_{ANO}^h(p^l)) = (c^h - c^l)/(2\beta - 1)^2 > 0$. Furthermore, for $\beta = 0$, $p_{ANO}^h(p^l)$ coincides with the equal mark-up line. Together these findings imply that for all admissible (α, β) $p_{ANO}^h(p^l)$ lies to the left and above the equal mark-up line. Here again the arrows in Figure S1 originating from $p_{ANO}^h(p^l)$ and from Point Z indicate the comparative statics with respect to α (no impact), β (positive impact; in the sense that $p_{ANO}^h(p^l)$ increases if β increases).

The critical prices $p_{ANU}^h(p^l)$ and $p_{ANO}^h(p^l)$ determine the expert's behavior in Area A. Whenever $p^h \in [p_{ANU}^h(p^l), p_{ANO}^h(p^l)]$ the expert will provide the appropriate treatment independent of the need of the customer. If $p^h < p_{ANU}^h(p^l)$ the expert will always provide low quality independent of the customer's needs; if $p^h > p_{ANO}^h(p^l)$ the expert will always provide high quality independent of the customer's needs. Note that the width of the area in which the expert provides the appropriate treatment is increasing in β : $p_{ANO}^h(p^l) - p_{ANU}^h(p^l) = \beta v/(1 - 2\beta)$.

Summing up: In Area A the expert's provision behavior is independent of α . Furthermore, an expert with $\beta > 0$ has a corridor around the equal mark-up line, where she provides the appropriate treatment, above this corridor she will overtreat, below the corridor she will undertreat. The width of the corridor is increasing in β .

In **Area B** (green line and blue dashed line) an expert providing appropriate treatment is better off than the customer if the latter needs t^h but she

is worse off if he needs t^l .

(BNU) The expert will not undertreat iff $p^h - c^h - \beta(2p^h - c^h - v) \geq p^l - c^l - \beta(2p^l - c^l)$, i.e. the condition is identical to (ANU): $p_{BNU}^h(p^l) = p_{ANU}^h(p^l)$. Thus the blue dashed line in Figure S1 (and S2) applies to both areas above \bar{p}^h .

(BNO) The expert will not overtreat (green solid line) iff: $p^l - c^l - \alpha(c^l + v - 2p^l) \geq p^h - c^h - \beta(2p^l - c^l - v)$ or

$$p^h \leq p_{BNO}^h(p^l) = \frac{(1+2\alpha)}{(1-2\beta)}p^l + \frac{1}{1-2\beta}[c^h(1-\beta) - c^l(1+\alpha) - v(\alpha+\beta)].$$

The comparative statics are $\frac{d}{d\alpha}(p_{BNO}^h(p^l)) = \frac{1}{2\beta-1}(v+c^l-2p^l) < 0$ and $\frac{d}{d\beta}(p_{BNO}^h(p^l)) = -\frac{1}{(2\beta-1)^2}(v+2v\alpha+2c^l\alpha-4p^l\alpha-c^h+2c^l-2p^l) > 0$. The first inequality (the one in the comparative statics with respect to α) follows from the fact that in Area B we have $p^l < \bar{p}^l = (v+c^l)/2$ and from $\beta < \frac{1}{2}$. The second inequality follows from the fact that at Point X in figures S1 and S2, defined as the point where $p_{BNO}^h(p^l) = \bar{p}^h$ we have $p^l = \frac{(v+c^l)}{2} - \frac{(c^h-c^l)}{2+4\alpha}$ independently of β while Point Z in figures S1 and S2, defined as the point where $p_{ANO}^h(\bar{p}^l) = p_{BNO}^h(\bar{p}^l) = \frac{v+c^l}{2} + \frac{(1-\beta)(c^h-c^l)}{1-2\beta}$, is increasing in β : $\frac{d}{d\beta}(p_{ANO}^h(\bar{p}^l)) = \frac{c^h-c^l}{(-1+2\beta)^2} > 0$ (see also the discussion of Points X, Y, and Z below).

Is $p_{BNU}^h(p^l)$ everywhere to the right and below of $p_{BNO}^h(p^l)$? At the right boundary of Area B (where $p^l = \bar{p}^l$) this is necessarily the case as $p_{ANO}^h(\bar{p}^l) = p_{BNO}^h(\bar{p}^l)$ and we know from the analysis of Area A that $p_{ANO}^h(p^l) > p_{ANU}^h(p^l) \forall p^l$. At the lower boundary (where $p^h = \bar{p}^h$) this is the case if and only if $(p_{BNO}^h)^{-1}(\bar{p}^h) < (p_{BNU}^h)^{-1}(\bar{p}^h)$ which is equivalent to

$$\alpha < \hat{\alpha} = \beta \frac{v - (c^h - c^l)}{(c^h - c^l) - 2v\beta}.$$

Note that this is the condition given in Footnote 12 of the main article.

Summing up: If $\alpha < \hat{\alpha}$ then provision behavior of the expert in Area B is qualitatively as shown in Figure S1, if $\alpha > \hat{\alpha}$ then it is as shown in Figure S2. In the former case a triangle limited by \bar{p}^h, \bar{p}^l , and $p_{BNO}^h(p^l)$ marks the area where the expert always provides the appropriate treatment and in the rest

of the area (above and to the left) the expert always provides t^h . In the latter case Area B contains regions where the expert always provides appropriate treatment (small triangle starting from \bar{p}^l and pointing downward), overtreats (large upper left region), undertreats (lower right region) and always provides the wrong treatment (small triangle sitting at \bar{p}^h and pointing upwards).

In **Area C** (red lines) the seller is worse off than the customer if she provides the appropriate treatment, independently of whether the customer needs t^h or t^l .

(CNU) The expert will not undertreat (red dashed line) iff: $p^h - c^h - \alpha(c^h + v - 2p^h) \geq p^l - c^l - \beta(2p^l - c^l)$ or

$$p^h \geq p_{CNU}^h(p^l) = \frac{(1 - 2\beta)}{(1 + 2\alpha)}p^l + \frac{1}{1 + 2\alpha}[c^h(1 + \alpha) - c^l(1 - \beta) + \alpha v].$$

The comparative statics of this condition with respect to α and β yield: $\frac{d}{d\alpha}(p_{CNU}^h(p^l)) = \frac{1}{(2\alpha+1)^2}(v - 2c^l\beta + 4p^l\beta - c^h + 2c^l - 2p^l) > 0$ and $\frac{d}{d\beta}(p_{CNU}^h(p^l)) = -\frac{1}{2\alpha+1}(2p^l - c^l) < 0$.

(CNO) The expert will not overtreat (red solid line) iff: $p^l - c^l - \alpha(c^l + v - 2p^l) \geq p^h - c^h - \alpha(c^h + v - 2p^h)$ or

$$p^h \leq p_{CNO}^h(p^l) = p^l + \frac{(1 + \alpha)(c^h - c^l)}{1 + 2\alpha}$$

This critical value does not depend on β . The comparative statics of this condition with respect to α yields $\frac{d}{d\alpha}(p_{CNO}^h(p^l)) = -\frac{c^h - c^l}{(2\alpha+1)^2} < 0$.

Summing up: If $\alpha < \hat{\alpha}$ then provision behavior of the expert in Area C is qualitatively as given in Figure S1, if $\alpha > \hat{\alpha}$ then it is qualitatively as in Figure S2. In the former case Area C contains regions where the expert always provides appropriate treatment (triangle starting from \bar{p}^h and pointing downward), overtreats (upper left region), undertreats (lower right region) and always provides the wrong treatment (triangle starting from the y-axis and pointing to the upper right boundary of the area). In the latter case the 'always appropriate treatment' region disappears from Area C.

In **Area D** (brown line and dashed red line) the expert who provides appropriate treatment is better off than her customer if he needs t^l but she is worse off if the customer needs t^h .

(DNU) The expert will not undertreat (red dashed line) iff: $p^h - c^h - \alpha(c^h + v - 2p^h) \geq p^l - c^l - \beta(2p^l - c^l)$, i.e. the condition is identical to the condition in paragraph (CNU). Thus, the expert will not undertreat if $p^h \geq p_{DNU}^h(p^l) = p_{CNU}^h(p^l)$.

(DNO) The expert will not overtreat iff $p^l - c^l - \beta(2p^l - c^l - v) \geq p^h - c^h - \alpha(v + c^h - 2p^l)$ or

$$p^h \leq p_{DNO}^h(p^l) = \frac{(1 - 2\beta)}{(1 + 2\alpha)}p^l + \frac{1}{1 + 2\alpha}[c^h(1 + \alpha) - c^l(1 - \beta) + (\alpha + \beta)v]$$

Since $p_{DNO}^h(\bar{p}^l) = \bar{p}^h + \frac{1}{2} \frac{c^h - c^l}{1 + 2\alpha} > \bar{p}^h$ and $\frac{d}{dp^l}(p_{DNO}^h(p^l)) > 0$ this condition is never binding, the respective line does not enter Area D.

Summing up: If $\alpha < \hat{\alpha}$ then provision behavior of the expert in Area D is qualitatively as given in Figure S1 (appropriate treatment in a small subarea to the left, undertreatment in the rest of the area), if $\alpha > \hat{\alpha}$ then it is qualitatively as in Figure S2 (undertreatment everywhere). Overtreatment is never an issue in Area D.

3 Characterizing the Intersections of the Boundary-Conditions

In the following we discuss some further characteristics of the diagram that hold with Fehr and Schmidt preferences independent of the parameters. For this purpose we define three points in the (p^l, p^h) space: the intersection of $p_{BNO}^h(p^l)$ and $p_{CNO}^h(p^l)$, the intersection of $p_{ANU}^h(p^l)$ and $p_{CNU}^h(p^l)$, and the intersection of $p_{ANO}^h(p^l)$ and $p_{BNO}^h(p^l)$. See figures S1 and S2 for an illustration.

Point X: First, we show that the intersection of $p_{BNO}^h(p^l)$ and $p_{CNO}^h(p^l)$ (Point X) always lies on \bar{p}^h . Then we study how changes in α and β affect the location of this intersection. To see the first part, define \tilde{p}^l as the p^l such that $p_{BNO}^h(\tilde{p}^l) = p_{CNO}^h(\tilde{p}^l)$. Then \tilde{p}^l is given as $\tilde{p}^l = (p_{BNO}^h)^{-1}(\bar{p}^h) = \frac{(v+c^l)}{2} - \frac{(c^h-c^l)}{(2+4\alpha)}$ and $p_{BNO}^h(\tilde{p}^l) = p_{CNO}^h(\tilde{p}^l) = \frac{v+c^h}{2}$. To see the second part, note $(p_{BNO}^h)^{-1}(\bar{p}^h)$

is independent of β . To study the change of this point in α , first remember that for $\alpha = 0$ the point lies on the equal mark-up line. Furthermore, $\frac{d}{d\alpha}(\hat{p}^l) = \frac{c^h - c^l}{(1+2\alpha)^2} > 0$. Thus, Point X moves to the right when α increases. In Figure S1 this is indicated by the red arrow originating in Point X.

Point Y: We show that the intersection of $p_{ANU}^h(p^l) = p_{BNU}^h(p^l)$ and $p_{CNU}^h(p^l)$ (Point Y) always lies on \bar{p}^h . Then we study how changes in α and β affect its location. Then we confirm the condition for Point Y lying to the right of Point X found earlier. To see the first part, define \tilde{p}^l as the p^l such that $p_{ANU}^h(\tilde{p}^l) = p_{CNU}^h(\tilde{p}^l)$. Then $\tilde{p}^l = (p_{ANU}^h)^{-1}(\bar{p}^h) = \frac{1}{2} \frac{v+(1-2\beta)c^l - (c^h - c^l)}{1-2\beta}$ and $p_{ANU}^h(\tilde{p}^l) = p_{CNU}^h(\tilde{p}^l) = \frac{v+c^h}{2}$. To see the second part, note that $(p_{BNO}^h)^{-1}(\bar{p}^h)$ is independent of α . The change of this point in β is given by $\frac{d}{d\beta}(\tilde{p}^l) = \frac{v-(c^h - c^l)}{(-1+2\beta)^2} > 0$. Thus, Point Y moves to the right when β increases. In Figure S1 this is indicated by the black arrow originating in Point Y. Third, note that Point Y lies to the right of Point X iff $\tilde{p}^l > \hat{p}^l$, which is equivalent to $\alpha < \hat{\alpha} = \beta \frac{v-(c^h - c^l)}{(c^h - c^l) - 2v\beta}$.

Point Z: We show that the intersection of $p_{ANO}^h(p^l)$ and $p_{BNO}^h(p^l)$ (Point Z) always occurs at \bar{p}^l and that the intersection necessarily lies above \bar{p}^h . Then we show that it moves upward when β increases and remains unchanged when α changes. First, solving $p_{ANO}^h(p^l) = p_{BNO}^h(p^l)$ for p^l yields $p^l = \bar{p}^l$. Furthermore, $p_{ANO}^h(\bar{p}^l) = \bar{p}^l + \frac{1}{2} \frac{c^h - c^l}{1-2\beta}$. Thus, Point Z lies above \bar{p}^l , is independent of α and moves upwards in β : $\frac{d}{d\beta}(p_{ANO}^h(\bar{p}^l)) = \frac{c^h - c^l}{(2\beta - 1)^2} > 0$.

4 Implications for Provision Behavior

We summarize all our findings in figures S1 and S2. The red and black arrows indicate the comparative statics results for changes in α (red) and β (black).

If $\alpha = \beta = 0$ all lines collapse to the equal mark-up line and points X, Y and Z are all on this line (Point X and Point Y collapse into a single point in this degenerate case). X, Y lie at the intersection of the equal mark-up line with \bar{p}^h and Point Z at the intersection of the equal mark-up line with \bar{p}^l .

If α increases Point X moves to the right, while Y and Z remain unaffected.

By contrast, if β increases, Point X remains unaffected while Point Y moves rightward and Point Z upwards. As can easily be seen, increases in β lead to increases in the area where appropriate treatment is provided. By contrast, increases in α decrease the area where appropriate treatment is provided.

Preempting the discussion in the next section note that $\alpha > 0$ and $\beta = 0$ corresponds to the borderline case between "inequality aversion" and "competitive preferences" (where $\alpha \geq 0$ and $\beta \leq 0$): In this borderline case inequality aversion has only negative consequences for experts' provision behavior. By contrast, $\alpha = 0$ and $\beta > 0$ is the borderline case between "inequality aversion" and "efficiency loving" preferences (where $\alpha \leq 0$ and $\beta \geq 0$). In this borderline case inequality aversion has only positive consequences for provision behavior.

5 Using the Fehr and Schmidt Model to Analyze other Forms of Distributional Preferences

By lifting the parameter restriction $\alpha \geq \beta \geq 0$ in the Fehr and Schmidt model the impact of (linear versions of) other forms of distributional preferences (EL, CO, and IL, see the main article for definitions) on expert's provision behavior can also be represented using the same figure. The respective assumptions for the linear (Fehr and Schmidt - FS) versions of the respective types of distributional preferences on α and β are:

(FS-IA) Inequity Aversion: $\alpha \geq \beta \geq 0$ and $\beta \leq \frac{1}{2}$.

(FS-EL) Efficiency Loving: $\beta \geq -\alpha \geq 0$ and $\alpha \geq -\frac{1}{2}$ and $\beta \leq \frac{1}{2}$.

(FS-CO) Competitive (spiteful) Preference: $\alpha \geq -\beta \geq 0$.

(FS-IL) Inequity Loving: $-\beta \geq -\alpha \geq 0$ and $\alpha > -\frac{1}{2}$.

How does the behavior of a **FS-EL** expert in the (p^l, p^h) space look like? Figure S1 helps to answer this question. Starting from the equal mark-up line, Point X moves to the left, Point Y to the right and Point Z upward.

Using the color scheme of the respective areas in Figure 6 of the main article, the resulting figure has a green corridor (with appropriate treatment) around the equal mark-up line, a yellow area (overtreatment) above and to the left of the green area and a red area (undertreatment) below and to the right. Thus, the resulting figure looks qualitatively as Figure 5 in the main article, the only difference being that the boundaries of the green corridor can have kinks.

How does the behavior of a **FS-CO** expert in the (p^l, p^h) space look like? Starting from the equal mark-up line, Point X moves to the right, Point Y to the left and Point Z downward. Using the color scheme of the respective areas in Figure 6 of the main article, the resulting figure has a blue corridor (with always wrong treatment) around the equal mark-up line, a yellow area (overtreatment) above and to the left of the then blue area and a red area (undertreatment) below and to the right. Thus, the resulting figure looks qualitatively the same as Figure 5 in the main article except that the green area is replaced by a blue area. An immediate implication is that for competitive preferences there exists no price vector that leads to provision of the appropriate treatment.

Finally, how does the behavior of a **FS-IL** expert in the (p^l, p^h) space look like? Starting from the equal mark-up line, Point X moves to the left, Point Y to the right and Point Z upward. Using the color scheme of the respective areas in Figure 6 of the main article, the resulting figure looks qualitatively as Figure 6 with two exceptions, a) the green area in Figure 6 becomes blue while the blue area becomes green; and b) the single point where the green and the blue area meet is to the left of the equal mark-up line and not to the right as in Figure 6 of the main article.

Figures




Figure S1: This diagram assumes $\alpha < \beta(v - (c^h - c^l)) / ((c^h - c^l) - 2v\beta)$ and is qualitatively the same as Figure 6 in the main document. Red arrows show the comparative statics in α , black arrows show the comparative statics in β

Above the dashed lines the expert does not undertreat. Below solid lines she does not overtreat. If a point lies between the two lines and the dashed line is below the solid line, the expert provides the appropriate treatment. If a point lies between the two lines and the dashed line is above the solid line the expert provides t^h when t^l is needed and t^l when t^h is needed. If a point lies below both lines the expert will always provide low quality, independent of the customer's need. If a point lies above both lines the expert will always provide high quality, independent of the customer's need.




Figure S2: The case with $\alpha > \beta(v - (c^h - c^l)) / ((c^h - c^l) - 2v\beta)$ - using the colour scheme of the main article, the blue area extends into area B.