

How Distributional Preferences Shape Incentives on (Experimental) Markets for Credence Goods

Rudolf Kerschbamer, Matthias Sutter and Uwe Dulleck*

This Version: May 2013

Abstract

Credence goods markets suffer from inefficiencies caused by superior information of sellers about the surplus-maximizing quality. While standard theory predicts that equal mark-up prices solve the credence goods problem if customers can verify the quality received, experimental evidence indicates the opposite. We identify a lack of robustness of institutional design with respect to heterogeneity in distributional preferences as a possible cause and design new experiments that allow for parsimonious identification of sellers' distributional types. Our results indicate that less than a fourth of the subjects behave according to standard theory's assumption, the rest behaving either in line with non-standard selfish or in accordance with non-trivial other-regarding preferences. We discuss consequences of our findings for institutional design and agent selection.

JEL Classifications: C72, C91, D82

Keywords: Credence Goods, Expert Services, Distributional Preferences, Other-Regarding Preferences, Behavioral Economics, Experimental Economics

*Kerschbamer: Department of Economics, University of Innsbruck (e-mail: rudolf.kerschbamer@uibk.ac.at); Sutter: Department of Public Economics, University of Innsbruck, and Department of Economics, University of Gothenburg (e-mail: matthias.sutter@uibk.ac.at); Dulleck: School of Economics and Finance, Queensland University of Technology (e-mail: uwe.dulleck@qut.edu.au).

Acknowledgements: A previous version of this paper was circulated as "The Impact of Distributional Preferences on (Experimental) Markets for Credence Goods" (IZA Discussion Paper 4647). We gratefully acknowledge financial support from the Austrian Science Fund (FWF; grant numbers P20796 and P22669), the Austrian National Bank (OeNB Anniversary Fund; grant number 13602) and the Australian Research Council (ARC; grant number DP110103653). Part of this research was conducted while the first author was visiting the Queensland University of Technology and he thanks its School of Economics and Finance for its great hospitality.

1 Introduction

A central topic in the field of information economics is the design of institutions or contracts that mitigate market inefficiencies resulting from the presence of asymmetric information. Almost all contributions to the literature build on the assumption of common knowledge that agents are rational own-money maximizers who behave as desired when kept indifferent in own-money terms – see Bolton and Dewatripont (2005) for a textbook coverage of this approach. In this paper we argue that while this assumption is harmless in some applications – because it results in institutions that are almost optimal if preferences are almost as assumed – it is misleading in others.

Specifically, we study markets for credence goods where inefficiencies result from superior information of sellers about the optimal quality for consumers, and where theory based on the aforementioned standard assumption predicts that equal mark-up prices solve the problem if customers can verify the quality received. This prediction is refuted by existing experimental evidence which indicates that markets with verifiability perform no better than markets without. We identify a lack of robustness of institutional design with respect to heterogeneity in distributional preferences as a possible cause and design new experiments that allow for parsimonious identification of sellers' distributional types based on their provision behavior. The results obtained in the implementation of the innovative experimental design indicate that less than a fourth of the subjects behave according to standard theory's assumption, with the rest behaving either in line with non-standard selfish or in accordance with non-trivial other-regarding preferences. These results support our explanation for the failure of verifiability to increase efficiency and suggest the search for an institutional design that is robust against preference heterogeneity as an important area for future research. Such research seems especially important for markets for credence goods where inefficient institutions potentially cause huge economic costs.¹

Credence goods markets are characterized by informational asymmetries between expert sellers and customers because customers are unable to iden-

¹Economically important credence goods markets include the market for medical care and that for car repair services. For the former the data in the WHO World Health Statistics (2009) implies that health care expenditures account for approximately 15% of GDP in the U.S. and are still rising. For the latter the online site researchandmarkets.com reports annual revenues of about \$ 90 billion for the U.S. auto repair industry, of which 70% originate from mechanical repair.

tify the quality they need, whereas expert sellers are able to do so (Darby and Karni, 1973). Typical examples are health care services, where the doctor is better informed than the patient on the disease the latter has and on the treatment he needs; car repair services, where the mechanic knows more about the type of service the vehicle needs than the owner; and taxicab rides in an unknown city, where the driver is better informed about the shortest route to the destination than the tourist. A second informational problem in markets for credence goods arises when the customer is unable to observe and verify the quality of service he received. For instance, in the market for medical treatments a patient might be unable to distinguish a cheap from an expensive drug infusion; and in the car repair market the owner might be unable to observe whether a broken part has been repaired or replaced.

The informational asymmetries on credence goods markets may cause a variety of problems and inefficiencies. Expert sellers may provide unnecessarily high quality (a case referred to as "overtreatment"), or insufficiently low quality ("undertreatment"), or they may charge for a higher quality than provided ("overcharging"). Such cases are not only a theoretical possibility, but are well documented in the literature – for instance, for the market for medical treatments (see, e.g. Hughes and Yule, 1992, Gruber and Owings, 1996, Gruber, Kim and Mayzlin, 1999, or Iizuka, 2007), for the one for car repairs (e.g., Wolinsky, 1993, Hubbard, 1998, or Schneider, 2012), and for that for taxi rides (c.f. Balafoutas et al., 2012).

An important finding in the theoretical literature is that verifiability ensures efficiency on markets for credence goods.² Verifiability applies if consumers are able to observe and verify the quality they receive so that expert sellers cannot charge for a quality that has not been provided.³ If verifiability applies, experts are predicted to choose equal mark-up prices. Under the mentioned standard assumption on preferences such prices induce them to provide the appropriate quality of the credence good. As a consequence,

²See Emons (1997, 2001), Pesendorfer and Wolinsky (2003), Alger and Salanié (2006), and Dulleck and Kerschbamer (2009) for research articles on the role of verifiability and Dulleck and Kerschbamer (2006) for a unifying model and a survey of the literature.

³Verifiability is likely to hold in many important credence goods markets. In some cases – for example pest control, equipment repair and dental services – the customer is present during provision and can ensure that services charged for have indeed been provided. In other cases verifiability is secured indirectly through the provision of appropriate evidence. For instance, in the market for car repairs, 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.

consumers – inferring experts’ incentives from posted prices – are predicted to interact and maximal efficiency of the market is expected.

We reanalyze experimental data in Dulleck, Kerschbamer and Sutter (2011) indicating that – contrary to theoretical prediction – verifiability fails to promote efficiency on credence goods markets. Indeed, the relative frequencies of market interaction, undertreatment and overtreatment do not differ significantly between two experimental treatments that are identical except that verifiability applies in one, but not in the other. The observed aggregate performance in both treatments is better in terms of efficiency than the standard prediction for a market without verifiability, but considerably worse than the prediction for a market with verifiability. These findings raise two questions whose answers are important for the understanding of – and the optimal design of institutions for – credence goods markets: Why is the performance of credence goods markets so poor in the presence of verifiability when all theoretical approaches predict verifiability to ensure efficiency? And why do markets without verifiability perform so much better than predicted?

In this paper, we argue that heterogeneity in the distributional preferences of credence goods sellers can explain both, why markets with verifiability perform so badly and why markets without verifiability perform substantially better than predicted. Key to our argument are the following observations: First, the standard solution to the credence goods problem for the case where the quality of the goods is verifiable – equal mark-up prices – is robust against positive deviations from the standard assumption on preferences, but non-robust against negative ones. Most importantly, it is non-robust against spiteful experts’ tendency to over- or undertreat customers and against inequality averse agents’ incentive to hurt the trading partner in the domain of disadvantageous inequality. Second, for the prediction for markets without verifiability the opposite is true – it is robust against negative deviations from the standard assumption on preferences, but non-robust against positive deviations. Most importantly, it is non-robust against efficiency loving experts’ tendency to provide the appropriate quality despite material incentives for undertreatment and against inequality averse agents’ positive attitude towards the trading partner in the domain of advantageous inequality. Moreover, there is no ‘cheap’ repair for the equal mark-up solution to the credence goods problem in the sense that it does not help to give up just (a small) part of the surplus to make it robust against (small) deviations from the standard model of preferences.

To receive empirical support for our explanation for the failure of ver-

ifiability to secure efficiency on markets for credence goods we design new experiments intended to identify an expert's distributional preferences from his provision behavior. Our main theoretical innovation here is the derivation of an experimental design that allows to identify a seller's distributional archetype without making any structural assumption on sellers' utility or motivation function meant to represent preferences. This distinguishes our approach from most of the rest of the literature on identification of type and intensity of distributional concerns which uses identification procedures that rely on strong structural assumptions regarding the form of the utility or motivational function – such as linearity, piecewise linearity, or specific forms of convexity, for instance.^{4,5}

We then implement our test for distributional concerns in new credence goods markets experiments. Our main findings are that (i) only a minority (of less than a fourth) of subjects behave according to the standard assumption of lexicographic maximization of first the own and then the other's material payoff; (ii) the behavior of a sizeable minority of subjects is consistent with other forms of selfish preferences; (iii) the behavior of a large majority of sellers is consistent with either a taste for efficiency (in the spirit of Andreoni and Miller, 2002, or Charness and Rabin, 2002) or inequality aversion (in the tradition of Fehr and Schmidt, 1999, or Bolton and Ockenfels, 2000); and (iv) a minority of subjects behaves spitefully or competitive (à la Levine, 1998, or Charness and Rabin, 2002). Hence, our empirical findings provide strong support for heterogeneity in distributional preferences and therewith for our explanation for the surprisingly low level of efficiency on credence goods markets in the presence of – and the surprisingly high efficiency level in the absence of – verifiability.

⁴For instance, the ring-test – originally developed by social psychologists to assess "social value orientations" and recently used by economists to identify type and intensity of distributional concerns (see, for instance, Offerman et al., 1996, Brosig, 2002, or Brandts et al., 2009) – is based on the assumption of *linear* preferences; the studies by Iriberry and Rey-Biel (2010), Cabrales et al. (2010) and Blanco et al. (2011) employ identification procedures based on the *piecewise linear* model originally introduced by Fehr and Schmidt (1999) as a description of self-centered inequality aversion and later extended by Charness and Rabin (2002) to allow for other forms of distributional concerns; and Andreoni and Miller (2002), Fisman et al. (2007) and Cox and Sadiraj (2012) check consistency with – and estimate parameters of – standard or modified *CES* utility functions.

⁵An exception is Kerschbamer (2011) who develops a test for distributional preferences that shares many features with the one proposed here. We discuss the relationship further in Section 3.

The remainder of the paper is organized as follows. Section 2 first introduces the basic setup of a credence goods market on which our experiments are based. It then presents predictions based on the mentioned standard assumption on preferences and the results from two experimental treatments in Dulleck et al. (2011) which only differ with respect to the presence or absence of verifiability.⁶ Section 3 presents our explanation for the low level of efficiency on credence goods markets in the presence and the high level of efficiency on markets in the absence of verifiability. The explanation is based on the presence of heterogeneity of distributional preferences, but neither on the exact fractions of different archetypes of distributional concerns nor on the exact shape of the preferences of the archetypes. Section 4 contains our main theoretical innovation – the design of a non-parametric test for distributional concerns that does not rely on any structural assumptions on preferences and that is completely nested in a market for credence goods. Section 5 implements the proposed experimental design and presents the results. Section 6 concludes with a discussion of our results and their implications for institutional design and for agent selection. Our main message there is that good performance of credence goods markets requires either robust institutions (combined with arbitrary experts) or ”good” experts (combined with arbitrary institutions).

2 Verifiability in Credence Goods Markets: Model, Standard Predictions and Experimental Evidence

This section is subdivided in several subsections. Subsection 2.1 introduces the simple model of a credence goods market on which our experiments are based, and Subsection 2.2 the experimental parameterization of the basic

⁶Dulleck et al. (2011) have a total of 16 experimental treatments (on the role of liability, verifiability, competition and reputation) of which we discuss only two here (and add two new ones). An earlier version of Dulleck et al. (2013) studies with data reported in the current paper a completely different question, namely whether a seller’s price-posting behavior is indicative of her intentions regarding provision and charging behavior. The actual version of that paper is based on new experiments with a different institutional framework and a different parameter constellation. Huck et al. (2007, 2010, 2012) have interesting experiments on the effect of prices and opportunities to build up a reputation on the performance of markets for experience (rather than credence) goods.

model used in Dulleck et al. (2011) – which will also be used in our new experiments reported in Section 5. Subsection 2.3 presents predictions on the role of verifiability based on the mentioned standard assumption on preferences, and Subsection 2.4 uses data from two experimental treatments in Dulleck et al. (2011) to show that verifiability does not have the predicted effects and that non-trivial distributional concerns seem to influence the behavior of some subjects.

2.1 Basic Model

Consumers are ex ante identical. They need a high quality, q^1 , of a particular (credence) good with probability h , and a low quality, q^0 , with probability $1-h$. Each consumer (he) is randomly matched with one seller (she) who sets prices p^1 and p^0 for the high, respectively low, quality (with $p^1 \geq p^0$). The seller has costs c^1 (c^0 , respectively) for the high (low) quality, with $c^1 > c^0$.

The consumer only knows the prices for the different qualities, but not the quality 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 quality the customer needs. Then she provides one of the two qualities and charges one of the two prices.

Customers in need of the low quality are sufficiently treated in any case (both if the seller chooses q^0 and if she chooses q^1). However, if the customer needs the high quality, then only q^1 is sufficient. A sufficient quality yields a value $v > 0$ for the customer, an insufficient quality 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 the quality received minus the price to be paid, whereas the seller receives as a monetary payoff the price charged minus the costs of the quality provided. More formally, let $\theta \in \{0, 1\}$ be the index of a customer's need in terms of quality, $\mu \in \{0, 1\}$ the index of the quality provided, and $\kappa \in \{0, 1\}$ the index of the quality charged for. Then the material payoff of the seller under price-vector (p^0, p^1) is

$$\pi_s(p^0, p^1, \mu, \kappa) = p^\kappa - c^\mu \tag{1}$$

while the customer receives

$$\pi_c(p^0, p^1, \theta, \mu, \kappa) = v - p^\kappa \text{ if } \theta \leq \mu \text{ and } -p^\kappa \text{ otherwise.} \quad (2)$$

<INSERT FIGURE 1 about here>

Figure 1 presents this game. Note, that this simple game captures all the idiosyncratic problems of credence goods markets discussed in the introduction. If a customer needs q^1 and the seller provides q^0 , we have *undertreatment*; if the customer needs q^0 and the seller provides q^1 , we have *overtreatment*; and if the seller charges p^1 when q^0 is provided, we have *overcharging*.

2.2 Experimental Design

In the following we introduce the experimental parameterization of the basic model used in Dulleck et al. (2011) which will also be used in our new experiments. We compare two treatments, one without verifiability (treatment *N-Endo*), and one with verifiability (treatment *V-Endo*).⁷ Treatment *N-Endo* corresponds to the game shown in Figure 1. Implementing verifiability means that consumers are able to observe and verify *ex post* the quality of the provided good (without knowing, however, whether this quality is the appropriate one). Therefore, in treatment *V-Endo* the last stage in Figure 1 is degenerate because the expert has to charge the price for the provided quality. Hence, with verifiability overcharging (and undercharging) are precluded, while over- and undertreatment are still possible.

In both treatments the customer's probability of needing the high quality is $h = 0.5$, and the value of a sufficient quality is $v = 10$. The costs of providing the low, respectively high, quality are $c^0 = 2$, and $c^1 = 6$. The prices posted by the sellers, p^0 and p^1 (with $p^0 \leq p^1$), have to be chosen in integer numbers from the interval $\{1, \dots, 11\}$. The outside option if no trade takes place between the seller and the customer is set to $o = 1.6$.

⁷The main difference in experimental design between the new experiments and those in Dulleck et al. (2011) is our reliance on (carefully designed) exogenously given prices for different qualities of the good rather than letting sellers endogenously decide on prices. To emphasize this difference we refer to the treatments *B/N* and *B/V* in Dulleck et al. (2011) as treatments *N-Endo* and *V-Endo* here, while the new treatments have names ending in *-Exo*.

Matching groups of eight subjects each were implemented, with four subjects as customers and four subjects as sellers. Role assignment was random at the beginning and fixed for all 16 periods in the experiment. In order to prevent attempts to build up a reputation as a reliable seller, there was random re-matching of customers and sellers within each matching group after each period. 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 treatments *N-Endo* and *V-Endo*.

2.3 Standard Prediction for the Role of Verifiability

Prediction 1 (Standard Prediction for the Role of Verifiability)

Under the assumption that subjects have standard preferences, in treatment N-Endo no interaction will take place, yielding no efficiency gains in the market. By contrast, in treatment V-Endo the expert will post $p^0 = 6$ and $p^1 = 10$ and the consumer will choose to enter the market and he will get the appropriate quality, yielding full efficiency in the market.

The following considerations lead to this prediction: The standard assumption of common knowledge that all agents are rational, risk-neutral and exclusively interested in their own material payoff implies that in treatment *N-Endo* the expert will always charge the higher price p^1 and always provide the cheaper quality q^0 . Anticipating this a consumer will only accept if $p^1 \leq (1 - h)v - o = 3.4$. But with such a p^1 even a defrauding seller earns less than the value of her outside option (because $(1 - h)v - c^0 < 2o$). In treatment *V-Endo* the expert cannot charge for a quality other than the provided one, and the provided quality depends on the mark-up $p^\mu - c^\mu$, $\mu \in \{0, 1\}$: An equal mark-up price-vector is defined as one that satisfies $p^1 - c^1 = p^0 - c^0$ and under the mentioned standard assumption on preferences (that if held indifferent in own-money terms the expert will provide in the best interest of the customer) it is predicted to induce provision of appropriate quality. An undertreatment (overtreatment, respectively) price-vector satisfies $p^1 - c^1 < p^0 - c^0$ ($p^1 - c^1 > p^0 - c^0$, respectively) and is predicted to induce provision of low (high) quality independently of the customer's need. Figure 2 shows in the space of price-vectors the set of equal mark-up price-vectors (those vectors lie on the green line with slope 1), the set of undertreatment price-vectors (red area below the equal mark-up line) and the set of overtreatment vectors (yellow area above the equal mark-up line).

Anticipating how an expert’s provision behavior depends on price-vectors, a consumer will accept an equal mark-up vector iff $p^1 \leq 10$, an undertreatment vector iff $p^0 \leq 3$, and an overtreatment vector iff $p^1 \leq 8$. Thus, to maximize profits, the expert will post the equal mark-up vector $(p^0, p^1) = (6, 10)$, which will be accepted by an own-money-maximizing, risk-neutral consumer.

<INSERT FIGURE 2 about here>

2.4 Experimental Results

Observation 1 (Experimental Results for the Role of Verifiability)

Compared to treatment N-Endo, verifiability has no significant impact on the frequency of interaction, the undertreatment rate, the overtreatment rate and overall efficiency. The overall performance in both treatments is better than the standard prediction for treatment N-Endo, but worse than the one for treatment V-Endo.

Table 1, Figure 3 and Figure 4 support this observation, leading us to reject both parts of Prediction 1: Contrary to the prediction efficiency gains and interaction rates are not significantly different between the two treatments and they are significantly higher than 0 and significantly lower than 1 in both.

<INSERT TABLES 1 and 2, as well as FIGURES 3 and 4 about here>

A possible explanation for the relatively high interaction rate and the relatively low undertreatment rate in *N-Endo* is experts having a taste for efficiency. Another possible explanation is that experts care for equitable payoffs. Support for the latter hypothesis comes from the analysis of price-posting behavior. Contrary to the theoretical prediction equal mark-up prices are very rare in *V-Endo* – they are chosen in less than 5% of all transactions. Table 2 reports the frequencies of the five most popular price-vectors posted by sellers in the two treatments. It is interesting to note that in treatment *V-Endo* only one equal mark-up vector is among the top 5 price-vectors, but it is not the predicted one. In both treatments the price-vector (6, 8) is by far the most frequently posted price-vector. If the seller always provided the appropriate quality and charged for it, then this price-vector would split the

gains from trade equally between the consumer and the seller in each of the two states (i.e. when the consumer needs the low, and when he needs the high quality). The prominence of this price-vector therefore suggests that a concern for relative payoffs plays a role for aggregate behavior in the experiment. Of course, these observations provide only a rough indication that distributional preferences may shape sellers' behavior. In Section 4 we are going to develop a simple parsimonious test for distributional preferences within the framework of a credence goods market which is then implemented in new experiments in Section 5. Before doing so we argue (in Section 3) that heterogeneity in distributional preferences can explain both, the (relative to the standard prediction) bad performance of markets with verifiability and the good performance of markets without verifiability.

3 Heterogeneity in Distributional Preferences and Robustness of Institutions

Key to our explanation for the relatively bad performance of credence goods markets with verifiability and the relatively good performance of markets without verifiability is the observation that the standard solution to the credence goods problem for the case where the quality of the goods is verifiable – equal mark-up prices – is robust against the positive impact of distributional preferences, but non-robust against the negative one. For the prediction for markets without verifiability the opposite is true – i.e., it is robust against negative deviations from the standard model of preferences, but not against positive deviations. It is important to note that our arguments here and in the next section will not depend on any assumptions on the distribution of different other-regarding preference-types in the population and on how exactly the preferences of the subjects look like. Specifically, our discussion here and in the subsequent section relies on the assumption that (experimental) credence goods sellers are heterogeneous and that their preferences can be represented by a utility or motivation function $U(\pi_s, \pi_c)$ satisfying the following three conditions:

- $\partial U / \partial \pi_s > 0$;
- $\text{sign}(\partial U / \partial \pi_c)$ depends (only) on whether $\pi_s \geq \pi_c$, or $\pi_s < \pi_c$; and
- $\partial U / \partial \pi_s > \partial U / \partial \pi_c$.

The first of those conditions is innocuous, it requires only that – holding the material payoff of the customer constant – the seller’s utility increases in own material payoff. This assumption is satisfied by all empirically relevant distributional preference types discussed in the economic literature. The second assumption is both permissive and restrictive, depending on the perspective taken. It is permissive because it allows for all major types of distributional preferences that have been discussed in the economics literature.⁸ The second assumption may be considered as restrictive because it implies (i) that preferences only depend on outcomes, not on the way they are achieved (this is the defining feature of distributional preferences); and (ii) that the reference point for the evaluation of allocations (if one is used) is an equal-material-payoffs allocation.⁹ The third assumption is fairly innocent 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 somewhat restrictive for allocations with $\pi_s > \pi_c$; its main purpose is to get a unique ”switching point” in the test proposed below, though, and it can be relaxed without changing results qualitatively. We call an expert:

- **SE - selfish** iff a) $\partial U/\partial\pi_c = 0$ for $\pi_s \geq \pi_c$; & b) $\partial U/\partial\pi_c = 0$ for

⁸This includes altruism (à la Andreoni and Miller, 2002) and surplus maximization (as discussed by Engelmann and Strobel, 2004, for instance); inequality aversion (as modelled, e.g., by Fehr and Schmidt, 1999, or Bolton and Ockenfels, 2000), difference aversion (Charness and Rabin, 2002) and egalitarian motives (Dawes et al., 2007; Fehr et al., 2008); maximin (Engelmann and Strobel, 2004), Rawlsian (Charness and Rabin, 2002) and Leontief preferences (Andreoni and Miller, 2002; Fisman et al., 2007); spiteful preferences (Levine, 1998) and concerns for relative income (Duesenberry, 1949); envy (Bolton, 1991; Kirchsteiger, 1994; Mui, 1995); and equity aversion (Charness and Rabin, 2002; Fershtman et al., 2012).

⁹Implication (i) excludes, among others, seller behavior that is influenced by reciprocity considerations (as modeled by Rabin, 2002, Dufwenberg and Kirchsteiger, 2004, and Falk and Fischbacher, 2006, for instance), or by beliefs on the payoff expectations of the customer (as in the guilt aversion models of Charness and Dufwenberg, 2006, and Battigalli and Dufwenberg, 2007). Implication (ii) is somewhat restrictive (for instance, there might exist sellers who consider it fair to get 10% more than their customer but unfair to get 20% more), but is still less restrictive than existing models of reference-dependent distributional concerns (as, for instance, those proposed by Bolton, 1991, Mui, 1995, Fehr and Schmidt, 1999, Bolton and Ockenfels, 2000, or Charness and Rabin, 2002) which *assume* that preferences change qualitatively at equality, while we only *allow* preferences to change qualitatively at equality. See Kerschbamer (2011) for a more detailed discussion on this assumption.

$$\pi_s < \pi_c;$$

- **EL - efficiency loving** iff a) $\partial U/\partial\pi_c > 0$ for $\pi_s \geq \pi_c$; & b) $\partial U/\partial\pi_c > 0$ for $\pi_s < \pi_c$;
- **SP - spiteful** iff a) $\partial U/\partial\pi_c < 0$ for $\pi_s \geq \pi_c$; & 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 $\pi_s \geq \pi_c$; b) $\partial U/\partial\pi_c \leq 0$ for $\pi_s < \pi_c$; & c) at least one of the two derivatives is different from 0;
- **IL - inequality loving** iff a) $\partial U/\partial\pi_c \leq 0$ for $\pi_s \geq \pi_c$; b) $\partial U/\partial\pi_c \geq 0$ for $\pi_s < \pi_c$; & c) at least one of the two derivatives is different from 0.

A selfish (**SE**) seller is a homo oeconomicus according to standard theory; she simply maximizes her own material payoff. 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 spiteful (**SP**) 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. 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 inequality loving (**IL**) expert is willing to sacrifice own material payoff to increase the difference between the payoffs of the two trading partners. Figure 5 displays typical indifference curves for the five archetypes of distributional preferences in the own-payoff/other-payoff space. It is important to note that our classification does not assume that the intensity of other-regarding concerns remains constant within a given domain – as, for instance, the piecewise linear Charness and Rabin (2002) model does.

<INSERT FIGURE 5 about here>

What can we say about the market behavior of credence goods sellers exhibiting those types of distributional concerns? Consider *markets without verifiability* (N -markets) first. For such markets the standard prediction – undertreatment and overcharging under each price-vector – is already a worst case scenario that leaves no room for deterioration. To see this, consider a **SP** expert, for instance. By always providing the cheaper quality q^0 and

always charging the higher price p^1 such an expert maximizes her material payoff while at the same time minimizing the payoff of the customer. Thus, she behaves exactly like a **SE** expert in N -markets. The same is true for other experts with negative attitudes towards customers – most importantly for **IA** experts in the domain of disadvantageous inequality. However, the positive side of distributional preferences potentially manifests itself in a better market outcome than predicted under standard preferences. To see this consider an **EL** expert who finds out that the customer needs high quality. By undertreating the customer instead of providing q^1 she increases her material payoff by $c^1 - c^0$ at a cost of $v > c^1 - c^0$ to the customer. Thus, if the additional profit the seller receives from providing q^0 instead of q^1 (i.e., $c^1 - c^0$) is small compared to the loss arising from undertreatment (i.e., v), and if the weight on π_c in her utility function is sufficiently high relative to the weight on π_s , she will refrain from undertreatment. The same is true for **IA** experts in the domain of advantageous inequality and for **IL** experts in the domain of disadvantageous inequality. In sum, in N -markets experts with negative distributional concerns behave exactly like experts with standard preferences while experts with positive distributional concerns tend to behave better than predicted by standard theory.

For the standard solution for *markets with verifiability* (V -markets), by contrast, we get the opposite result. To see this, note that the standard prediction for equal mark-up prices – appropriate treatment independent of the level of the mark-up – is already a best-case scenario that leaves no room for improvement. However, negative deviations from the standard assumption of preferences easily manifest themselves in the market outcome because hurting the customer involves no cost under equal mark-up prices. To see this, consider an **EL** expert, for instance. Since the material payoff of the customer enters positively in her utility function, she will act in the interest of the consumer along the equal mark-up line, where helping the customer involves no cost. Furthermore, since $\partial U / \partial \pi_c > 0$ in both domains (i.e., in the domain of advantageous inequality and in the domain of disadvantageous inequality) the **EL** expert will provide the appropriate quality even under price-vectors that deviate (slightly) from the equal mark-up rule. Thus, **EL** experts necessarily provide appropriate quality in a corridor along the equal mark-up line – as shown in Figure 6 – but they do not perform better

than **SE** experts at the equal mark-up line.¹⁰ The same is true for other experts with positive attitudes towards customers – most importantly for **IA** experts in the domain of advantageous inequality. However, the negative side of distributional preferences easily manifests itself in a worse market outcome than predicted under standard preferences. To see this, consider a **SP** expert, for instance. Since the material payoff of the customer enters negatively in her utility function, she necessarily provides high quality to a consumer who needs low quality and low quality to a consumer who needs high quality along the equal mark-up line where hurting the customer involves no cost. Furthermore, since $\partial U/\partial \pi_c < 0$ in both domains the **SP** expert will always provide the wrong quality even under price-vectors that deviate (slightly) from the equal mark-up rule. The same is true for other experts with negative attitudes towards customers – most importantly for **IA** experts in the domain of disadvantageous inequality. Together these observations do not only explain the poor performance of equal mark-up prices in *V-Endo*, they also explain why equal mark-up prices are very rarely chosen in *V-Endo*. More importantly, there is no cheap repair for this problem in the sense that there is simply no price-vector that induces an **SP** expert, for instance, to provide the appropriate quality in a *V*-market. Her provision behavior is rather (qualitatively) like the one shown in Figure 6 with the important exception that she will necessarily always provide the wrong (instead of the appropriate) quality in a corridor along the equal mark-up line.

<INSERT FIGURE 6 about here>

4 Identification of Distributional Preferences in Markets for Credence Goods

The discussion in the previous section assumes that there is heterogeneity in distributional preferences in the (experimental) expert population. The challenge is, of course, to show that empirically. In this section we derive a parsimonious test for the identification of distributional preferences in the framework of a credence goods market. The test relies on exactly the same primitive assumptions on preferences as the discussion in the previous section.

¹⁰Point Ω and the other price-vectors indicated by bullet points in Figure 6 are not important for the arguments in this section – we will refer to them in the next section.

Our starting point in deriving the test is the observation that in the space of possible price-vectors there is exactly one (and only one) that allows for a neat discrimination between the above defined preference types from the provision behavior in a credence goods market with verifiability. Looking at Figure 6 it is the price-vector referred to as 'Point Ω '. It is defined as follows:

Definition 1: *The price-vector $\Omega = (p_{\Omega}^0, p_{\Omega}^1)$ has $p_{\Omega}^0 = (v + c^1)/2 - (c^1 - c^0)$ and $p_{\Omega}^1 = (v + c^1)/2$.*

To discuss the properties of this price-vector we have first to define (and discuss the location of) the three dashed lines in Figure 6. The *upward sloping* dashed line is the equal mark-up line. It connects all price-vectors with $p^1 - p^0 = c^1 - c^0$ implying that the expert receives exactly the same material payoff independently of whether she provides q^0 or q^1 at points on this line. The *horizontal* dashed line connects all price-vectors where the expert and the customer receive exactly the same material payoff if the expert (correctly or incorrectly) provides q^1 . Thus, this line is defined by $\pi_s(p^0, p^1, \mu = 1, \kappa = 1) = \pi_c(p^0, p^1, \theta = 1, \mu = 1, \kappa = 1) = \pi_c(p^0, p^1, \theta = 0, \mu = 1, \kappa = 1) \iff p^1 = (v + c^1)/2$. And the *vertical* dashed line connects all price-vectors where the expert and the customer receive exactly the same material payoff if the expert correctly provides q^0 . Thus, this line is defined by $\pi_s(p^0, p^1, \mu = 0, \kappa = 0) = \pi_c(p^0, p^1, \theta = 0, \mu = 0, \kappa = 0) \iff p^0 = (v + c^0)/2$. Since Point Ω is at the intersection of the upward sloping and the horizontal dashed it has $p^0 = (v + c^0)/2 - (c^1 - c^0)/2$ implying that this point is necessarily to the left of the vertical dashed line – where we have $p^0 = (v + c^0)/2$.

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 needs the cheaper quality, q^0* . If the expert provides the appropriate quality, she induces a payoff allocation (π_s, π_c) with disadvantageous inequality. Why? Because Point Ω is strictly to the left of the vertical dashed line along which both parties get exactly the same material payoff if the expert correctly provides q^0 . If the expert provides the expensive quality instead, she induces an equal-material-payoff allocation, that is, an allocation with $\pi_s = \pi_c$ – this follows from the fact that Point Ω is on the horizontal dashed line. Furthermore, since Point Ω is on the equal mark-up line, the expert's own material payoff is the same in both allocations! What does this imply for provision behavior? An **EL** expert and an **IL** expert will necessarily decide for the asymmetric allocation (by providing the cheaper quality to a customer who needs the cheaper quality) because the own material payoff

is the same in both allocations while the customer's payoff is higher in the asymmetric than in the symmetric allocation (relevant for **EL**), or because disadvantageous inequality is present in the asymmetric but absent in the symmetric allocation (relevant for **IL**). By contrast, a **SP** and an **IA** expert necessarily decide for the symmetric allocation (by providing the expensive quality to a customer who needs the cheaper one) because the own material payoff is the same in both allocations while the customer's payoff is lower in the symmetric than in the asymmetric allocation (relevant for **SP**), or because disadvantageous inequality is absent in the symmetric but present in the asymmetric allocation (relevant for **IA**).

Now assume that the *customer needs the expensive quality*, q^1 . If the expert provides the appropriate quality q^1 , then she induces the equal-material-payoff allocation discussed in the previous paragraph. This follows from the fact that the material payoff of both parties for providing the expensive quality is independent of the quality need of the customer. If the expert provides q^0 instead, she induces a payoff allocation (π_s, π_c) with advantageous inequality. This follows from the fact that Point Ω has $p^0 = (v + c^0)/2 - (c^1 - c^0)/2$ which exceeds $c^0/2$ because $v > (c^1 - c^0)$. Furthermore, since Point Ω is on the equal mark-up line, the expert's own material payoff is the same in both allocations. Thus, an **EL** expert and an **IA** expert will necessarily decide for the symmetric allocation (by providing the expensive quality to a customer who needs the expensive quality) because the own material payoff is the same in both allocations while the customer's payoff is higher in the symmetric than in the asymmetric allocation (relevant for **EL**), or because advantageous inequality is absent in the symmetric but present in the asymmetric allocation (relevant for **IA**). By contrast, a **SP** and an **IL** expert necessarily decide for the asymmetric allocation (by providing the cheap quality to a customer who needs the expensive one) again because the own material payoff is the same in both allocations while the customer's payoff is lower in the asymmetric than in the symmetric allocation (relevant for **SP**), or because advantageous inequality is present in the asymmetric but absent in the symmetric allocation (relevant for **IL**).

In sum, if we observe the decision of an expert under the price-vector located at Point Ω in Figure 6 twice, once combined with the consumer needing the low quality and once combined with the consumer needing the high quality, then we can infer her distributional preference type with some precision. Specifically, calling the strategy of providing the appropriate quality in both cases '*appropriate treatment*', the strategy of providing the expensive quality

in both cases 'overtreatment', the strategy of providing the cheap quality in both cases 'undertreatment', and the strategy of providing the expensive quality when the cheap quality is needed and the cheap quality when the expensive one is needed 'always wrong treatment' we get the following result:

Proposition 1 (Impartial Distributional Preferences) *Consider the price-vector Ω as defined in Definition 1. Under this price-vector: a) appropriate treatment is consistent with SE and **EL** preferences but inconsistent with IA, SP and IL; b) overtreatment is consistent with SE and **IA preferences** but inconsistent with EL, SP and IL; c) undertreatment is consistent with SE and **IL** preferences but inconsistent with IA, SP and EL; d) always wrong treatment is consistent with SE and **SP** preferences but inconsistent with IA, EL and IL.*

Proof. Follows immediately from the text preceding the result. ■

Testing the provision behavior under the price-vector Ω is like eliciting impartial distributional preferences, because under this price-vector 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 here. Based on the predictions for Point Ω we now change p^0 slightly, while keeping p^1 constant, in order to test whether (experimental) sellers are willing to give up own material payoff to help or hurt the customer. Referring back to Figure 6 an increase in p^0 corresponds to a move along the horizontal dashed line to the right, while a decrease in p^0 corresponds to a move to the left. In terms of payoff allocations, moving to the right (left, respectively) of Point Ω in Figure 6 means that we increase (decrease) the expert's payoff for providing q^0 at the cost (for the benefit) of the payoff to the consumer, while keeping the payoffs for both parties for providing q^1 constant at the equal-material-payoff allocation $(\pi_s, \pi_c) = ((v - c^1)/2, (v - c^1)/2)$. 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, we get the following monotonicity result:

Lemma 1 (Monotonicity) *Consider two price-vectors, the price-vector Ω from Definition 1 and a second vector, Ψ , which has the same p^1 as Ω (i.e., $p_\Psi^1 = p_\Omega^1$) but a different p^0 (i.e., $p_\Omega^0 \neq p_\Psi^0$). If $p_\Omega^0 < p_\Psi^0$ ($p_\Omega^0 > p_\Psi^0$, respectively) then – keeping the consumer's need with respect to quality constant – an expert who provides q^0 (q^1 , respectively) under Ω must provide q^0 (q^1 , respectively)*

under Ψ .

Proof. First note that providing q^1 yields the equal-material-payoff allocation $\pi_s = \pi_c = (v - c^1)/2$ independently of the consumer's need and independently of whether Ω or Ψ is the relevant price-vector. By contrast, the payoff allocation from providing q^0 depends on both, the consumer's need and the type of contract. Suppose first the *consumer needs* q^0 . Then providing q^0 under Ω yields $\pi_s = (v - c^1)/2$ and $\pi_c = (v - c^1)/2 + (c^1 - c^0)$, while providing q^0 under Ψ yields $\pi_s = (v - c^1)/2 + \varepsilon$ and $\pi_c = (v - c^1)/2 + (c^1 - c^0) - \varepsilon$, where $\varepsilon > 0$ for $p_\Omega^0 < p_\Psi^0$ and $\varepsilon < 0$ for $p_\Omega^0 > p_\Psi^0$. Now suppose the *consumer needs* q^1 . Then providing q^0 under Ω yields $\pi_s = (v - c^1)/2$ and $\pi_c = (v - c^1)/2 + (c^1 - c^0) - v$, while providing q^0 under Ψ yields $\pi_s = (v - c^1)/2 + \varepsilon$ and $\pi_c = (v - c^1)/2 + (c^1 - c^0) - v - \varepsilon$, where $\varepsilon > 0$ for $p_\Omega^0 < p_\Psi^0$ and $\varepsilon < 0$ for $p_\Omega^0 > p_\Psi^0$. 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) . ■

Proposition 1 and Lemma 1 together imply:

Proposition 2 (Partial Distribution Preferences) *Consider the price-vectors Ω and Ψ from Lemma 1. Then observing* a) *appropriate treatment under Ω and Ψ is only consistent with **EL** preferences (but inconsistent with **SE**, **IA**, **SP** and **IL**); b) *overtreatment under Ω and overtreatment, appropriate treatment or always wrong treatment under Ψ with $p_\Omega^0 < p_\Psi^0$ is only consistent with **IA** preferences (but inconsistent with **SE**, **EL**, **SP** and **IL**); c) *undertreatment under Ω and undertreatment, appropriate treatment or always wrong treatment under Ψ with $p_\Omega^0 > p_\Psi^0$ is only consistent with **IL** preferences (but inconsistent with **SE**, **IA**, **SP** and **EL**); d) *always wrong treatment under Ω and always wrong under Ψ is only consistent with **SP** preferences (but inconsistent with **SE**, **IA**, **EL** and **IL**).****

To understand Proposition 2, the test to be applied in the next section (and the term 'partial' distribution preferences) consider an **IA** seller, for instance. From the arguments above we know that such an expert has to overtreat a customer under price-vector Ω . Increasing p^0 slightly while keeping p^1 constant creates a tension between a higher own monetary payoff and more inequality and vice versa for an **IA** expert. By deciding for overtreatment or switching to appropriate treatment (or always wrong treatment) she reveals a positive willingness to pay for reducing inequality because own-money-maximization would ask for undertreatment. The argument for sellers with other kinds of distributional preferences is similar.

5 Implementing the Test in Lab Experiments

5.1 Experimental Parameters and Procedures

To test for and classify the distributional preferences of sellers, we ran new experiments using a design based on the results derived in the previous section. The timing of the game was exactly the same as in the game described in Section 2, except for the first stage: Instead of letting sellers post their prices themselves, the price-vector in a given period was chosen (exogenously through the software) with equal probability from the set $\{(3,8), (4,8), (5,8), (6,8), (7,8)\}$. This set of vectors has two characteristics: First and foremost, it includes the equal mark-up vector Ω characterized in Proposition 1 – it is the vector $(4,8)$. Starting from this price-vector it then varies p^0 as described in Lemma 1 and Proposition 2. The allocations implied by the the equal mark-up vector Ω and by the other price-vectors in the set are displayed in Figure 7.¹¹ Second, this set of price-vectors includes the four most frequently chosen price-vectors in treatment *V-Endo* (see Table 2). We call the experimental treatment with this (exogenously given) set of price-vectors *V-Exo1*. In order to check whether the inclusion of the price-vector $(3,8)$ – which was not among the most frequently posted price-vectors in treatment *V-Endo* – has any impact on behavior, we also ran an experimental treatment where the exogenously determined price-vector was chosen with equal probability only from the four most frequently chosen price-vectors $(4,8)$, $(5,8)$, $(6,8)$, and $(7,8)$. We call this treatment *V-Exo2*. We ran four sessions with 16 subjects each both for *V-Exo1* and for *V-Exo2*. Hence, a total of 128 subjects

¹¹Providing the high quality induces the equal-material-payoff allocation $(2,2)$ independent of the needed quality under each price-vector in the set. If the *customer needs low quality*, the seller (implicitly) chooses between this allocation (by inefficiently providing high quality) and the allocation corresponding to the respective price-vector (as indicated in the figure) on the line (with slope -1) *above* the equal-material-payoff allocation (by efficiently providing low quality). If the *customer needs high quality*, the choice is between the equal-material-payoff point (by efficiently providing high quality) and the respective point on the line (with slope -1) *below* the equal-material-payoff allocation (by inefficiently providing low quality). Consider the price-vector $(4,8)$, for instance: For the case where the customer needs low quality the expert's choice between appropriately providing q^0 and inappropriately providing q^1 corresponds to a choice between $(\pi_s, \pi_c) = (2,2)$ and $(\pi_s, \pi_c) = (2,6)$, a point located vertically *above* $(2,2)$; if the customer needs high quality instead, then the expert's choice between appropriately providing q^1 and inappropriately providing q^0 corresponds to a choice between $(2,2)$ and $(2,-4)$, which is located vertically *below* $(2,-4)$.

participated in the new experiments (with no subject having participated in the experiment reported in Section 2). Sessions lasted less than 1.5 hours and average earnings were about 15 Euro.

<INSERT FIGURE 7 about here>

5.2 Experimental Results

Since we do not find any significant differences in behavior between *V-Exo1* and *V-Exo2* we report pooled data from both treatments. In the data analysis we first looked at violations of monotonicity according to Lemma 1. It turns out that 45 out of 64 sellers (70%) behave in line with the statement over all 16 periods of the experiment. Taking into account that some learning may go on in early periods, we decided to focus on the final 12 periods only (i.e., periods 5 to 16). In those periods the behavior of 56 out of 64 sellers (88%) respects the monotonicity condition. 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 Lemma 1, we had to exclude 3 from further analysis due to lack of data caused by customers' opting out.¹² Our data analysis is therefore based on 53 sellers.

Observation 2 (Identification of Distributional Preferences) *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 a fifth of sellers shows behavior that is consistent with strong inequality aversion – they over- or undertreat customers if this behavior reduces inequality (or turns disadvantageous into advantageous inequality) even if it also reduces their own monetary payoff. Adding up strong*

¹²Our criteria for inclusion/exclusion were as follows: We included all experts who had treated under price vector Ω at least one customer needing q^0 AND at least one customer needing q^1 . 50 of the 56 sellers were included under this rule. From the remaining 6 sellers, we included those where the data was consistent with exactly one of the distributional types introduced in Section 3. 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.

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.

Table 3 provides a summary of the data. To read it properly, note that sellers who are classified as either weak **EL**, weak **IA**, weak **SP**, or weak **IL** types are also classified as weak **SE** types. This has to be the case because weak **EL**, **IA**, **SP** and **IL** types behave exactly as the strong version of the respective type as 'impartial spectators' (that is, when there is no trade off between own material payoff and a fairness standard), i.e. at price-vector (4, 8) in Figure 7. Once p^0 varies, weak **EL**, **IA**, **SP** and **IL** types act exactly like (strong) **SE** types, because their own material payoff is at stake.¹³ Thus, for relative frequencies (given in parentheses in Table 3) to add up to 100%, one has to add up 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.

<INSERT TABLE 3 about here>

An important insight is that the behavior of only a minority of individuals is consistent with standard theory's assumption, i.e. that sellers always follow their monetary incentives and in case of indifference they act in the interest of customers. Less than a fourth of experimental sellers (those in the category "weak EL") exhibit behavior that is consistent with this assumption. This is an important insight for several reasons. First, it is important for the current application –institutional design for credence goods markets under verifiability– because it provides an explanation for both, why equal mark-up price-vectors do not work as predicted by theory, and why such vectors were not chosen in the endogenous pricing conditions of Dulleck et al. (2011). And secondly it is important for institutional design for markets plagued by asymmetric information more generally, because it suggests that institutional design based on the standard assumption on preferences might yield bad incentives for a majority of agents. The results reported in Table 3 also confirm the heterogeneity in distributional preferences on which our discussion in Section 3 was based – some sellers care for efficiency, some for

¹³Formally, 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 approaches zero". Note, however, that the limiting behavior is different for the four non-SE types!

‘equality of payoffs’, some do not care for the well-being of others (or for efficiency) at all.

Heterogeneity in preferences and behavior is a well established fact, of course. Indeed, it has been observed in many other games, for instance, in public goods games (Fischbacher, Gächter and Fehr, 2001, Fischbacher and Gächter, 2010) and in gift-exchange games (Fehr et al., 1997). Also, in the literature on identification of type and intensity of distributional concerns, heterogeneity is well known (see, e.g. Andreoni and Miller, 2002, Charness and Rabin 2002, Engelmann and Strobel, 2004, or Fisman et al., 2007). Our main contributions here are (i) that our identification procedure depends only on a small set of primitive assumptions on preferences, which is in contrast to much of the previous literatures; and (ii) that our test for distributional preferences is completely nested in a market for credence goods which might help to alleviate the concern that the results of elicitation procedures based on dictator games are not robust and not easy to extend to other important economic situations. Since our results are qualitatively in line with those reported by the path-breaking dictator game studies by Andreoni and Miller (2002), Charness and Rabin (2002), Engelmann and Strobel (2004) and Fisman et al. (2007) that reservation against such studies seems exaggerated, however.

6 Conclusions

This paper has argued that heterogeneity in distributional preferences provides an explanation for both, why credence goods markets with verifiability fail to reach efficient outcomes and why markets without verifiability perform considerably better than predicted by standard theory. Key to our argument is the observation that the standard prediction for markets without verifiability is non-robust against the presence of agents with positive distributional concerns, while the standard solution to the credence goods problem for the case where the quality of the goods is verifiable – equal mark-up prices – is non-robust against the presence of agents with negative distributional concerns.

To receive support for our explanation we have designed a parsimonious experiment that allows for clean discrimination between different preference types from their provision behavior. One important feature of our experimental design is that the discrimination does not depend on any structural

assumptions on the utility or motivational function meant to represent preferences. The experimental design rather directly tests the key characteristics of different variants of distributional preferences that have been discussed in the economic literature. A second important design feature is that our test for distributional concerns is completely nested in a market for credence goods.

While the design of a clean test for distributional concerns within a credence goods context is the key theoretical innovation of the present paper, important conclusions for credence goods-markets and, more generally, for markets with asymmetric information can be drawn from our experimental results. Specifically we found in the implementation of our experimental design that less than a fourth of the experimental sellers behave according to standard theory's assumption (that all agents are rational own-money maximizers who behave as desired if held indifferent in own-money terms), the rest behaving either in line with non-standard selfish or in accordance with non-trivial other-regarding preferences. An immediate implication is that institutional design based on the standard assumption of lexicographically maximizing agents yields bad incentives for a majority of agents. Another implication of our experimental results is that there are agents that behave appropriately independently of the institutional design. Taken together these two observations have two immediate consequences, one for institutional design, the other for agent selection.

Designing the Right Institutions: What is needed for a well-performing market is not a perfect institution for one type of agent, but rather an institution that is 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. (2011) have shown, 'liability' is a quite robust institution in markets for credence goods. '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, especially for markets for credence goods. Imposing liability, for instance, generates other problems or may be impossible to achieve.¹⁴ As a

¹⁴On the one hand, liability requires a form of verifiability of the outcome. Especially 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

consequence, selecting the "right" agents for jobs involving experts' services becomes particularly important. Instead of choosing doctors, mechanics or computer specialists exclusively according to their training, customers or their representatives should worry more about the attitudes of these experts towards their customers. Selecting the right agents may also help to solve problems created by uncertainty over input costs: 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 experts with the right distributional –or more generally speaking, social– preferences. Public policy might step in here, for instance, by screening candidates for crucial studies (as medicine, for instance) not only after their performance in entry exams but also in accordance with their social track record. Since the 'effort cost' for performing social activities is arguably lower for more 'consumer-friendly' types, a CV featuring an impressive track record of volunteer work might well act as a screening device.

(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 – a problem that has previously been discussed by Taylor (1995).

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] Andreoni, J. and Miller, J. (2002), Giving According to GARP: An Experimental Test of the Consistency of Preferences for Altruism, *Econometrica* **70**, 737-753.
- [3] Balafoutas, L., Beck, A., Kerschbamer, R. and Sutter, M. (2012), What Drives Taxi Drivers? A Field Study on Fraud in a Market for Expert Services, *Review of Economic Studies*, forthcoming.
- [4] Battigalli, P. and Dufwenberg, M. (2007), Guilt in Games, *The American Economic Review Papers and Proceedings* **97**, 170-176.
- [5] Blanco, M., Engelmann, D. and Normann, H. (2011), A Within-Subject Analysis of Other-Regarding Preferences, *Games and Economic Behavior* **72**, 321-338
- [6] Bolton, G. (1991). A Comparative Model of Bargaining: Theory and Evidence, *The American Economic Review* **81**, 1096-1136.
- [7] Bolton, G. E. and Ockenfels, A. (2000), ERC – A Theory of Equity, Reciprocity and Competition, *American Economic Review* **90**: 166-193.
- [8] Brandts, J., Riedl, A. and van Winden, F. (2009), Competitive Rivalry, Social Disposition, and Subjective Well-Being: An Experiment, *Journal of Public Economics* **93**, 1158-1167.
- [9] Brosig, J. (2002). Identifying Cooperative Behavior: Some Experimental Results in a Prisoner's Dilemma Game, *Journal of Economic Behavior and Organization* **47**, 275–290.
- [10] Cabrales, A., Miniaci, R., Piovesan, M. and Ponti, G. (2010), Social Preferences and Strategic Uncertainty: An Experiment on Markets and Contracts, *American Economic Review* **100**, 2261-2278.
- [11] Charness, G. and Dufwenberg, M. (2006), Promises and Partnership, *Econometrica* **74**, 1579-1601.

- [12] Charness, G. and Rabin, M. (2002), Understanding Social Preferences with Simple Tests, *Quarterly Journal of Economics*, **117**: 817-869.
- [13] Cox, J. and Sadiraj, V. (2012), Direct Tests of Individual Preferences for Efficiency and Equity, *Economic Inquiry* 50,920–931.
- [14] Darby, M. and Karni, E. (1973), Free Competition and the Optimal Amount of Fraud, *Journal of Law and Economics* **16**, 67-88.
- [15] Dawes, Ch., Fowler, J., Johnson, T., McElreath, R. and Smirnov, O. (2007), Egalitarian Motives in Humans, *Nature* **446**, 794-796.
- [16] Duesenberry, J. (1949), *Income, Saving and the Theory of Consumer Behavior*, Harvard University Press, Cambridge, MA.
- [17] Dufwenberg, M. and Kirchsteiger, G. (2004), A Theory of Sequential Reciprocity, *Games and Economic Behavior* **47**, 268-298.
- [18] Dulleck, U. and Kerschbamer, R. (2006), On Doctors, Mechanics, and Computer Specialists: The Economics of Credence Goods, *Journal of Economic Literature* 44: 5-42.
- [19] 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.
- [20] Dulleck, U., Kerschbamer, R. and Sutter, M. (2011), The Economics of Credence Goods: On the Role of Liability, Verifiability, Reputation and Competition, *American Economic Review* 101: 526–55.
- [21] Dulleck, U., Johnston, D., Kerschbamer, R, and Sutter, M. (2013), The Good, The Bad, and the Naive: Do Fair Prices Signal Good Types or Do They Induce Good Behaviour?, University of Innsbruck, mimeo.
- [22] Emons, W. (1997), Credence Goods and Fraudulent Experts, *Rand Journal of Economics* **28**, 107-119.
- [23] Emons, W. (2001), Credence Goods Monopolists, *International Journal of Industrial Organization* **19**, 375-89.

- [24] Engelmann, D. and Strobel, M. (2004), Inequality Aversion, Efficiency, and Maximin Preferences in Simple Distribution Experiments, *American Economic Review* **94**, 857-869.
- [25] Falk, A. and Fischbacher, U. (2006), A Theory of Reciprocity, *Games and Economic Behavior* **54**, 293-315.
- [26] Fehr, E., Bernhard, H., and Rockenbach, B. (2008), Egalitarianism in Young Children, *Nature* **454**, 1079-1083.
- [27] Fehr, E., Gächter, S. and Kirchsteiger, G. (1997), Reciprocity as a Contract Enforcement Device, *Econometrica* **65**: 833-860.
- [28] Fehr, E., Kirchsteiger, G. and Riedl, A. (1993), Does Fairness Prevent Market Clearing? An Experimental Investigation, *Quarterly Journal of Economics* **108**, 437-459.
- [29] Fehr, E. and Schmidt, K. (1999), A Theory of Fairness, Competition, and Cooperation, *Quarterly Journal of Economics* **114**: 817-868.
- [30] Fershtman, C., Gneezy, U. and List, J. (2012), Equity Aversion: Social Norms and the Desire to be Ahead, *American Economic Journal Micro* **4**, 131-144.
- [31] Fischbacher, U. (2007), Z-tree: Zürich Toolbox for Ready-Made Economic Experiments. *Experimental Economics* **10**: 171-178.
- [32] Fischbacher, U. and Gächter, S. (2009), Social Preferences, Beliefs, and the Dynamics of Free Riding in Public Good Experiments, *American Economic Review* **100**: 541-556.
- [33] Fisman, R., Kariv, S. and Markovits, D. (2007), Individual Preferences for Giving, *American Economic Review* **97**, 1858-1876.
- [34] Fischbacher, U., Gächter, S. and Fehr, E. (2001), Are People Conditionally Cooperative? Evidence from a Public Goods Experiment, *Economic Letters* **71**: 397-404.
- [35] 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.

- [36] 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.
- [37] Gruber, J. and Owings, M. (1996), Physician Financial Incentives and Cesarean Section Delivery, *Rand Journal of Economics* **27**, 99-123.
- [38] Hubbard, T. N. (1998), An Empirical Examination of Moral Hazard in the Vehicle Inspection Market, *RAND Journal of Economics* **29**, 406-426.
- [39] Huck, S., Lünser, G. and Tyran, J.-R. (2007), Pricing and Trust, UCL Working Paper.
- [40] Huck, S., Lünser, G. and Tyran, J.-R. (2010), Consumer networks and Firm Reputation: A First Experimental Investigation, *Economics Letters* **108**, 242-244.
- [41] Huck, S., Lünser, G. and Tyran, J.-R. (2012), Competition Fosters Trust, *Games and Economic Behavior* **76**, 195-209.
- [42] 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.
- [43] Iriberry, N. and Rey-Biel, P. (2009), Elicited Beliefs and Social Information in Modified Dictator Games: What do Dictator Believe Other Dictators Do?, Economics Working Papers 1137, Department of Economics and Business, Universitat Pompeu Fabra.
- [44] Iizuka, T. 2007. Experts' Agency Problems: Evidence From the Prescription Drug Market in Japan, *RAND Journal of Economics* **38**: 844-862.
- [45] Kerschbamer R. (2011), The Geometry of Distributional Preferences and a Non-Parametric Identification Approach, University of Innsbruck, mimeo.
- [46] Kirchsteiger, G. (1994), The Role of Envy in Ultimatum Games, *Journal of Economic Behavior and Organization* **25**, 373-89.
- [47] Levine, D. (1998), Modeling Altruism and Spitefulness in Experiments, *Review of Economic Dynamics* **1**, 593-622.

- [48] Mui, V. (1995), The Economics of Envy, *Journal of Economic Behavior and Organization* **26**, 311-336.
- [49] Offerman, T., Sonnemans, J. and Schram, A. (1996), Value Orientations, Expectations and Voluntary Contributions in Public Goods, *Economic Journal* **106**, 817-845.
- [50] 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.
- [51] Rabin, M. (1993). Incorporating Fairness into Game Theory and Economics, *The American Economic Review* **83**, 1281-1281.
- [52] Schneider, H. (2012), A Field Experiment to Measure Agency Problems in Auto Repair, *Journal of Industrial Economics* **60**, 406-433.
- [53] Sutter, M., Haigner, S. and Kocher, M. (2010), Choosing the Stick or the Carrot? - Endogenous Institutional Choice in Social Dilemma Situations, *Review of Economic Studies* **77**, 1540-1566.
- [54] Taylor, C. (1995), The Economics of Breakdowns, Checkups, and Cures, *Journal of Political Economy* **103**, 53-74.
- [55] Wolinsky, A. (1993), Competition in a Market for Informed Experts' Services, *Rand Journal of Economics* **24**, 380-398.
- [56] Wolinsky, A. (1995), Competition in Markets for Credence Goods, *Journal of Institutional and Theoretical Economics* **151**, 117-131.

7 Tables and Figures

Table 1: Summary Statistics for *N-Endo* and *V-Endo*

Averages per Period	<i>N-Endo</i> ^a	<i>V-Endo</i> ^a
Interaction	0.45	0.50
Efficiency ^b	0.18	0.16
Undertreatment ^c	0.53	0.60
Overtreatment ^d	0.06	0.05
Overcharging ^e	0.88	-
Profit Seller	2.69	2.58
Profit Customer	1.00	1.06

^a none of the variables is significantly different between the two treatments (using two-sided Mann-Whitney U-tests with matching groups of 8 subjects as independent observations).

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

^c customer needs q^1 , but seller provides q^0

^d customer needs q^0 , but seller provides q^1

^e seller provides q^0 , but charges p^1 (with $p^1 > p^0$ and customer needs q^0)

Table 2: Most Popular Price-Vectors in *N-Endo* and *V-Endo*

Treatment <i>N-Endo</i>			Treatment <i>V-Endo</i>		
(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: Classification of Individual Behavior in *V-Exo*

distributional type	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%)
SP (spiteful)	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%)

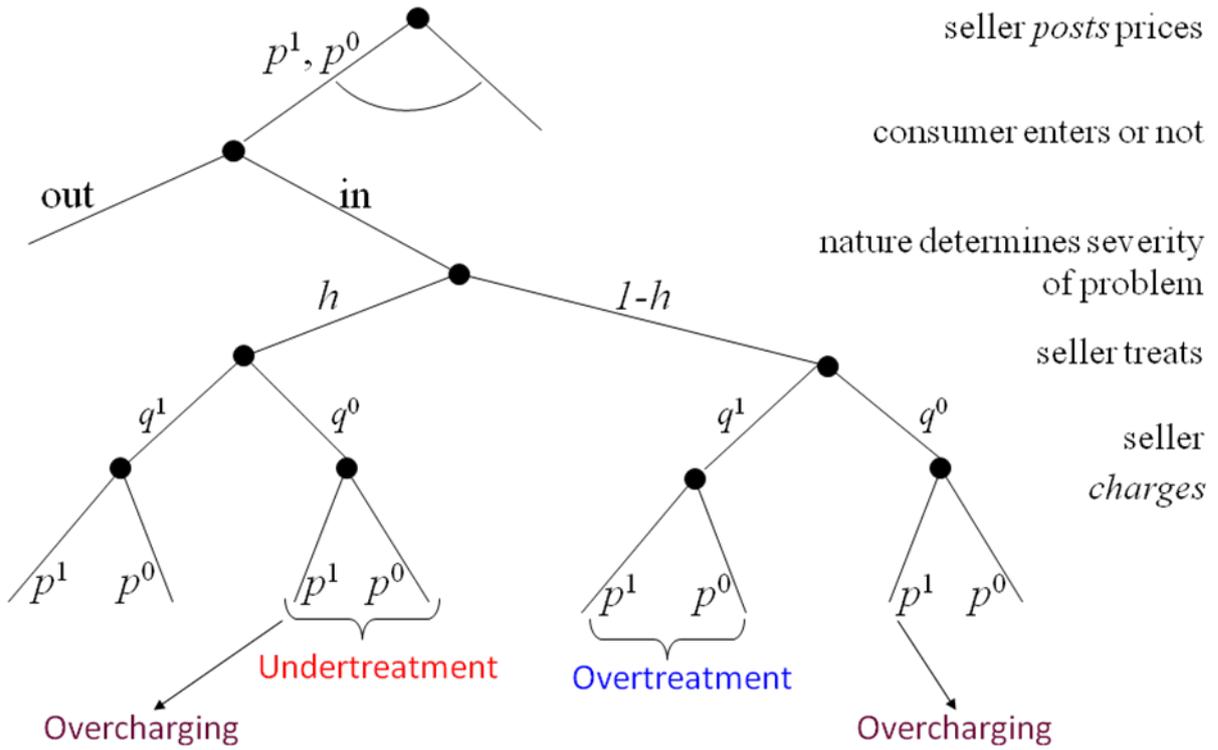


Figure 1: The Credence Goods Game

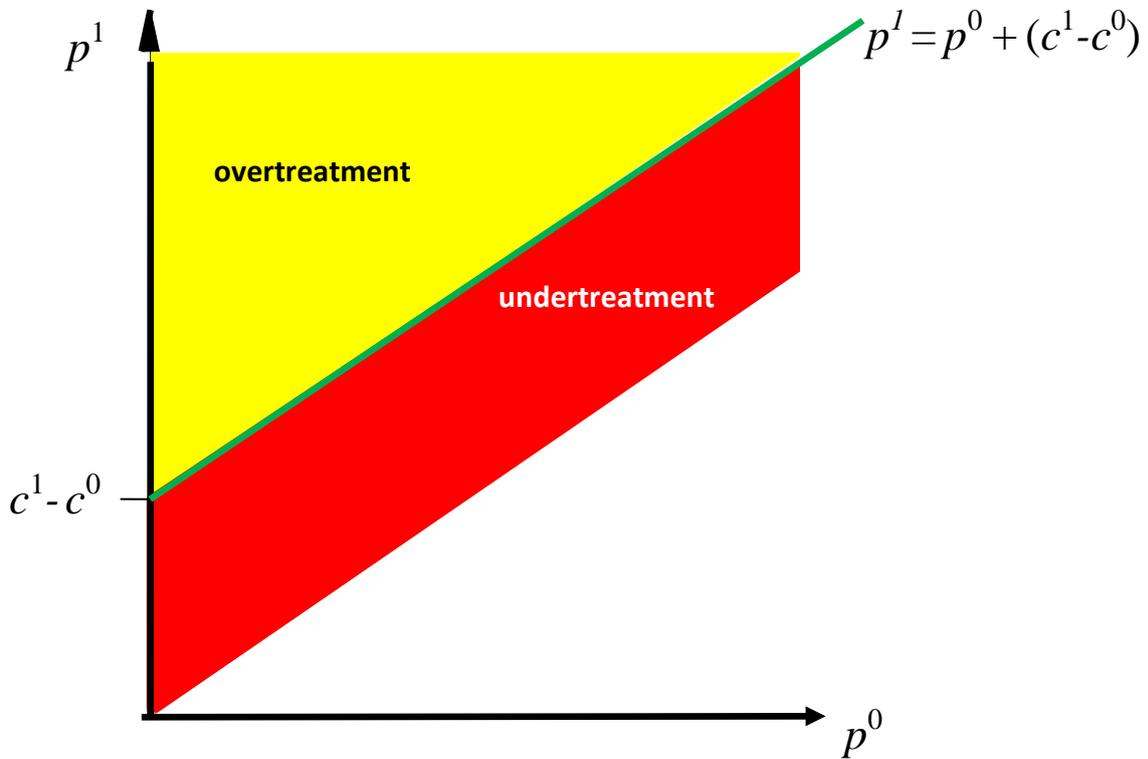


Figure 2: Standard Prediction for Provision Behavior under Verifiability

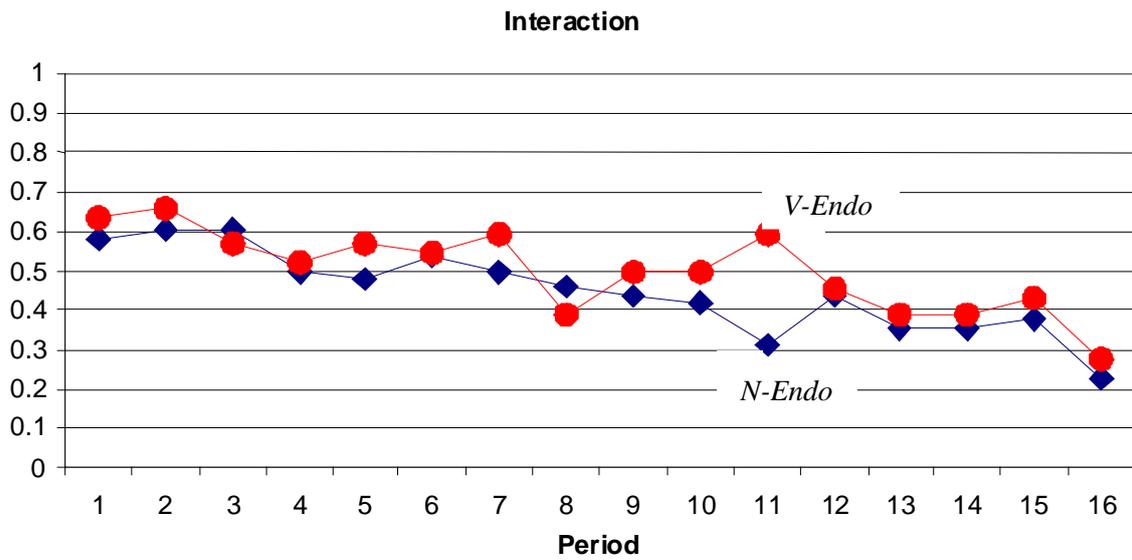


Figure 3: Relative Frequency of Interaction in *N-Endo* and *V-Endo*

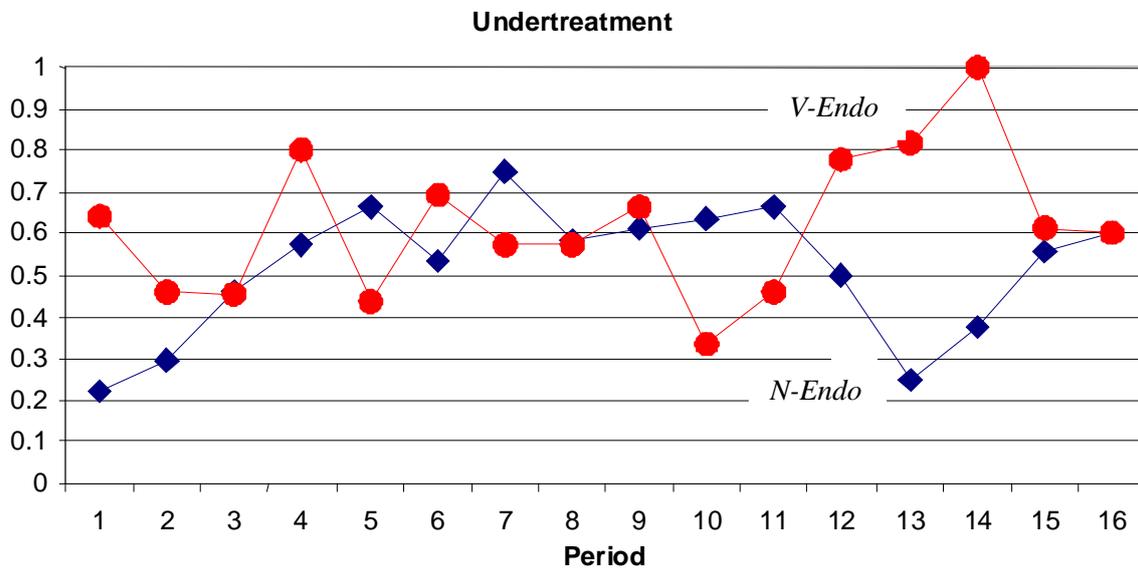


Figure 4: Relative Frequency of Undertreatment in *N-Endo* and *V-Endo*

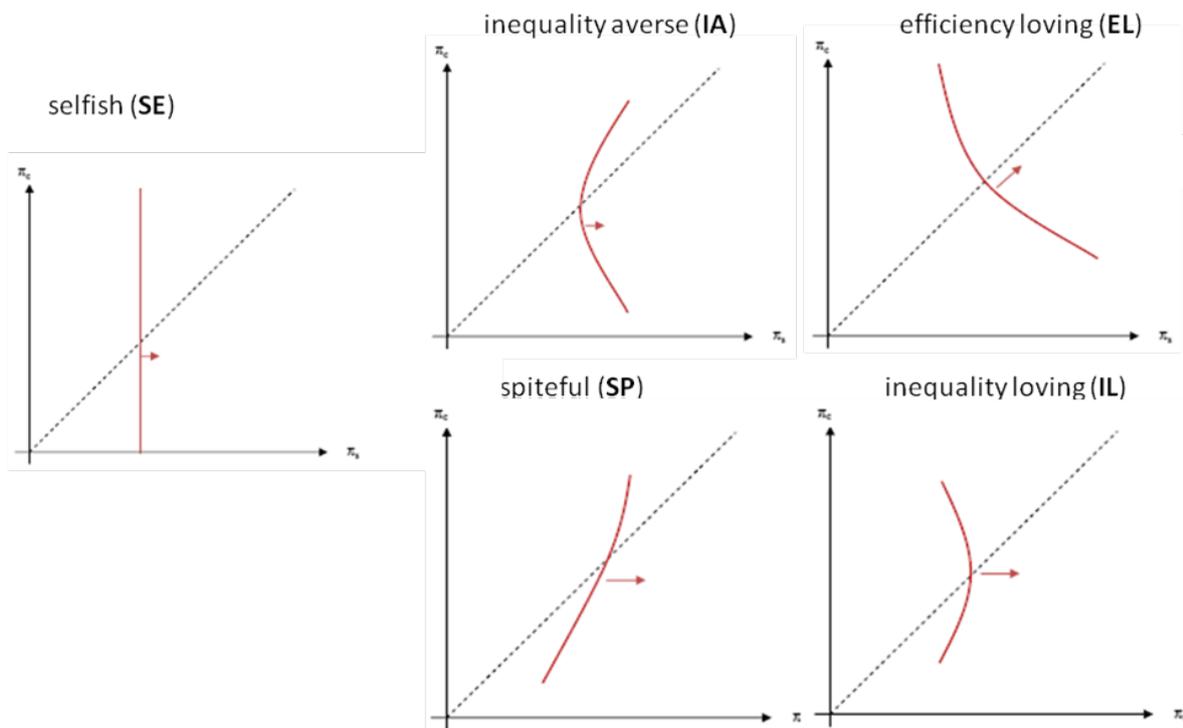


Figure 5: Indifference Curves of SE, IA, EL, SP and IL Experts in (π_s, π_c) Space

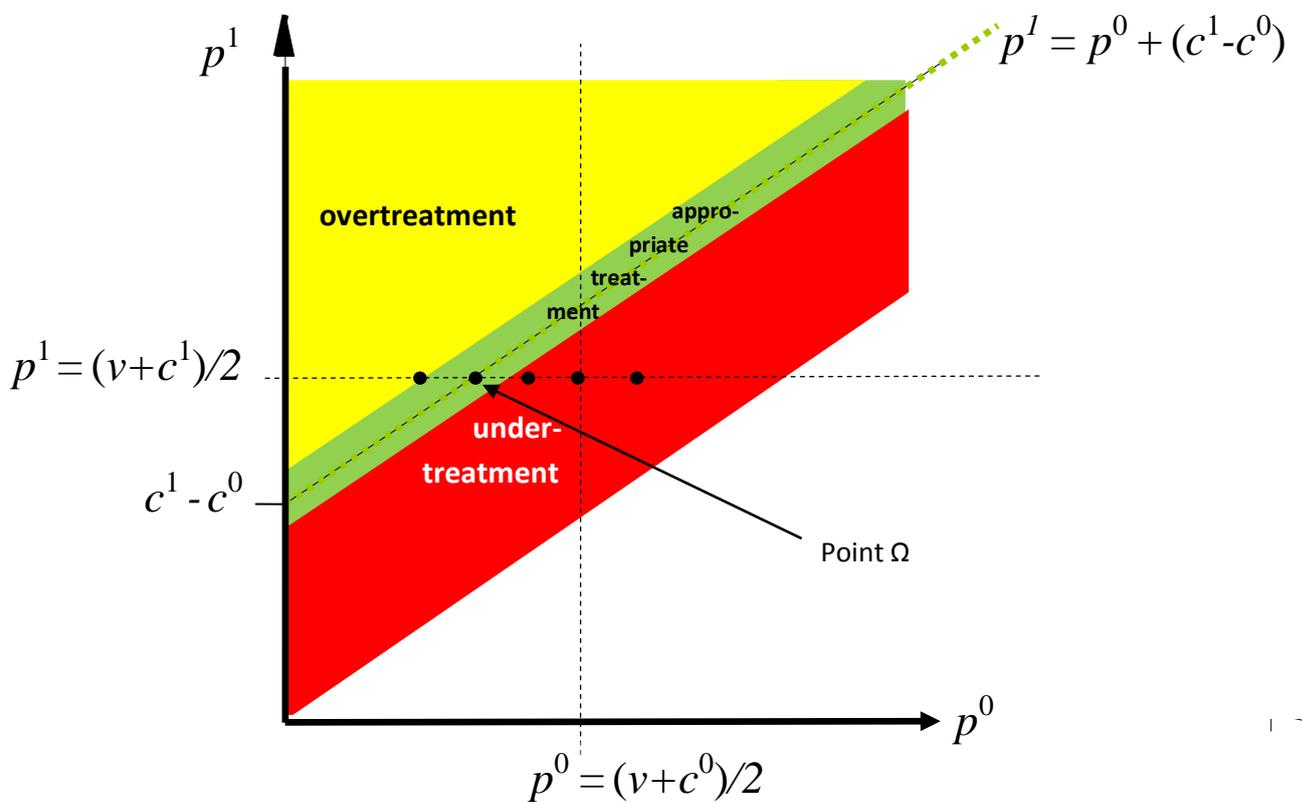


Figure 6: Provision Behavior of an EL Expert under Verifiability

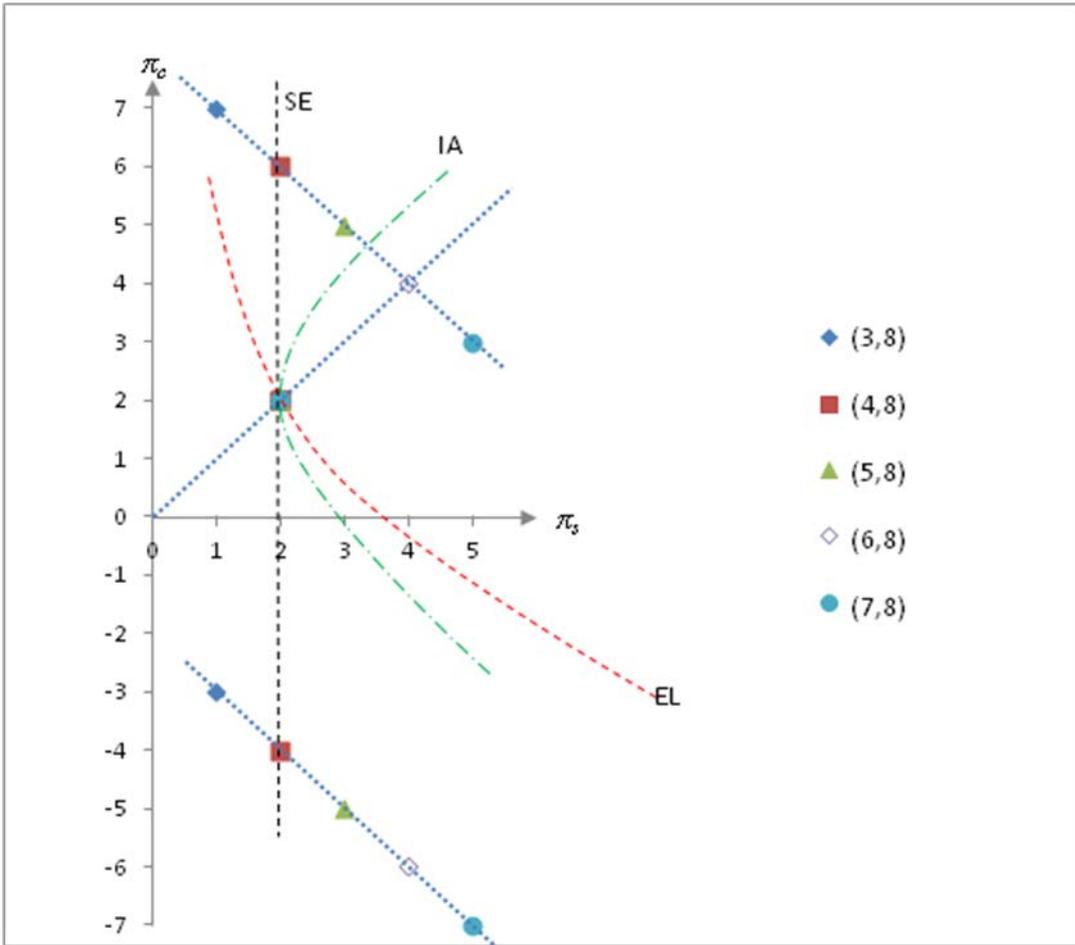


Figure 7: Possible Combinations of Buyer's and Seller's Material Payoffs (for different price-vectors and depending on whether the buyer needs q^0 or q^1)

Experimental Instructions for the *-Exo* Treatments (not intended for publication)

INSTRUCTIONS FOR THE EXPERIMENT

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

2 Roles and 16 Rounds

This experiment consists of **16 rounds**, each of which consists of the same sequence of decisions. This sequence of decisions is explained in detail below.

There are 2 kinds of roles in this experiment: **player A** and **player B**. At the beginning of the experiment you will be randomly assigned to one of these two roles. On the first screen of the experiment you will see which role you are assigned to. Your role remains the same throughout the experiment.

A player A interacts with a player B. This pair of players **changes** for each round. Therefore you are interacting in every round with a **new** player (of the other role).

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

Overview of the Sequence of Decisions in a Round

Each round consists of a maximum of 2 decisions which are made consecutively. Decision 1 is made by player B and decision 2 is made by player A. In each round 2 prices will be announced before players make their decisions. These prices are set for a given round. This price setting is referred to in the following as “Decision 0”.

Short Overview of the Sequence of Decisions in a Round

0. The prices for action I and action II are announced to both players.
1. Player B decides whether he/she wants to interact with player A. If he/she chooses No, the round ends.

If player B chooses to interact then

2. Player A (but **not** player B) is informed about the type of player B. There are two possible types of player B: he/she is of either type I or type II. Player A has to choose an action: either action I or action II. He/she then receives the price for the chosen action valid for this round. This price has to be paid by player B.

Detailed Illustration of the Decisions and Their Consequences Regarding Payoffs

Decision 0

In case of an interaction **player A** has to choose between two actions, action I and action II, in Decision 2. Each chosen action causes costs which are as follows:

Action I results in a **cost of 2 points** (=currency of the experiment) for player A.

Action II results in a **cost of 6 points** for player A.

Player A receives from player B the valid price for the action he/she chooses in Decision 2 if player B decides to interact with him/her. **At Decision 0 the valid prices for action I and action II for this round are announced to both players.**

Decision 1

Player B decides whether he/she wants to interact with player A.

If he/she wants to do so, then player A chooses an action in Decision 2 and he/she receives the valid price for this action from player B.

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

Decision 2

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

Player A gets to know the type of player B before he/she makes Decision 2. Then player A chooses an action, either action I or action II, and receives the corresponding price (valid for the respective round).

An **action** is **sufficient** under the following conditions:

- a) In case player B is a type I player and player A chooses either action I or action II.
- b) In case player B is a type II player and player A chooses action II.

An action is **not sufficient** if player B is of type II and player A chooses action I.

Player B receives 10 points, if the **action** chosen by player A is **sufficient**. **Player B** receives **0 points** if the **action** chosen by player A is **not sufficient**. In both cases player B has to pay the valid price for the chosen action.

At no time player B will be informed about whether he/she is of type I or a type II in any given round.

Payoffs

If player B chooses not to interact in Decision 1 (*decision “No” of player B*) then both players receive **1.6 points** for this round.

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

Player A receives the **price** (denoted in points, as announced in Decision 0) for the action chosen in Decision 2, **less the cost** of this action.

Player B's payoff depends on whether the Decision 2 of player A was sufficient or not.

- a) If the action was sufficient, **player B** gets **10 points less** the price for the action chosen by player A in Decision 2.
- b) If the action was not sufficient, **player B** has to pay the price for the action chosen by player A in Decision 2.

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

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

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