

Credence Goods, Costly Diagnosis, and Subjective Evaluation*

Helmut Bester[†]

Matthias Dahm[‡]

May 2015

Abstract

We study contracting in a credence goods model with two incentive problems. First, moral hazard occurs because the expert's choice of diagnosis effort is not observable. Second, payments can depend only on the consumer's subjective evaluation because only he observes the outcome of a treatment. We show that the first-best solution can *always* be implemented if the parties' discount rate is zero; an increase in the discount rate makes it more difficult to obtain the first-best. We further show that the first-best is also always implementable if separation of diagnosis and treatment is possible.

Keywords: credence goods, information acquisition, moral hazard, subjective evaluation, physician remuneration.

JEL Classification No.: D82, D83, D86, I11,

*We thank Vincent Anesi, Andreas Asseyer, Gianni De Fraja, Paula González, Peter Hammond, Matthias Lang, Inés Macho-Stadler, Marc Möller, Pau Olivella, Alex Possajennikov, Elke Renner, Daniel Seidmann, Silvia Sonderegger, Roland Strausz, Cecilia Testa, Justin Mattias Valasek and Ori Weisel for helpful discussions and suggestions. The usual disclaimer applies. Helmut Bester acknowledges support by the German Science Foundation (DFG) through SFB/TR 15. Matthias Dahm acknowledges the support of the Government of Spain under project ECO2010-19733.

[†]Freie Universität Berlin, School of Business and Economics, Boltzmannstr. 20, D-14195 Berlin (Germany); email: hbester@wiwiss.fu-berlin.de

[‡]University of Nottingham, School of Economics, University Park, Nottingham NG7 2RD (UK); email: Matthias.Dahm@nottingham.ac.uk

1 Introduction

This paper analyses contracting between a consumer (or principal) and an expert (or agent) in the presence of two incentive problems. First, there is a *moral hazard problem*. The consumer relies on the expert's advice in order to choose one of two services (or treatments). He can, however, not observe whether the expert exerted diagnostic effort and if so which signal he obtained. Second, there is a *problem of subjective evaluation*. As success and failure are only observed by the consumer and not publicly verified, payments for treatments can only rely on the principal's subjective evaluation, which might be misrepresented. In this paper we analyse the optimal design of contracts in such a situation and provide conditions under which the first-best can be reached.

The environment considered is the by now standard credence good problem.¹ In such a situation there is an important information asymmetry between the consumer and the expert. Even if the consumer can determine which service he received, he does not know whether he really needed an expensive high quality service or whether a less costly service of lower quality would have been sufficient. Consequently, incentives for opportunistic behaviour arise, because a self-interested expert who also provides the service might give inappropriate advice and remain undetected. Two types of inefficiencies can appear. On one hand, the quality of the treatment can be too low and not solve the principal's problem; we refer to this as *undertreatment*. On the other hand, *overtreatment* might occur, because when only low quality is needed, high quality is not valued higher than low quality.

These inefficiencies have potentially important implications. As many other countries, over recent years the U.S. experienced large increases in health care spending that many think are unsustainable. Dr. Donald M. Berwick, a former administrator of the Centers for Medicare and Medicaid Services, listed overtreatment as one of the key reasons for 'waste' in health care, saying that: "Much is done that does not help patients at all, and many physicians know it."² In fact, in Berwick and Hackbarth (2012, p. 1514) he writes that

¹The concept of credence goods was introduced by Darby and Karni (1973). Unlike experience goods, a credence good has important properties that the consumer cannot detect even after consumption. Classical examples include medical and legal advice, a variety of repair services, real estate services or taxi services. More recent applications of the concept of credence goods include consumer markets for food (Caswell and Mojdzuska (1996), Giannakas (2002)), the environmental and social impact of production (Feddersen and Gilligan (2001), Baron (2011)), technologically advanced consumption goods with many options (Dulleck and Kerschbamer (2009)), the newspaper industry (Gabszewicz and Resende (2012)), auditing services (Causholli et al. (2013), Knechel (2013)), financial services (Brown and Minor (2013)), contracting for infrastructure projects (Dulleck et al. (2013)), and fiscal restraints in public finance (Dulleck and Wigger (2015)).

²See e.g. Robert Pear, "Health Official Takes Parting Shot at 'Waste'," The New York Times 3, December 2011. For a discussion of the conceptual difficulties in identifying wasteful health care spending see Fuchs (2009). The ABIM Foundation (established by the American Board of Internal Medicine, an organization

overtreatment is “the waste that comes from subjecting patients to care that, according to sound science and the patients’ own preferences, cannot possibly help them—care rooted in outmoded habits, supply-driven behaviors, and ignoring science. Examples include excessive use of antibiotics, use of surgery when watchful waiting is better, and unwanted intensive care at the end of life for patients who prefer hospice and home care. We estimate that this category represented between \$ 158 billion and \$ 226 billion in wasteful spending in 2011.”³

It is well known that in the credence good problem the first–best can be reached, provided the expert can determine the agent’s type of problem without incurring any cost to himself (Dulleck and Kerschbamer (2006)). In order to do so the contract has to establish equal markup payments, which make the expert indifferent between treatments. Unequal markups would bias the expert towards one alternative, precluding truthful reporting. We depart from this credence good setting in three ways.

First, we depart from much of the credence goods literature by endowing the expert with a potentially less efficient diagnosis technology. First of all, we suppose that becoming informed requires exerting costly diagnosis effort. This creates a *moral hazard problem*, because whether effort has been exerted and if so which information has been obtained is not observable. As a result, equal markup contracts no longer provide effort incentives. The reason is that the expert gains the markup if he chooses based on prior information, while this markup is reduced by the effort cost when he exerts diagnostic effort. Moreover, we allow the expert’s signal to be noisy. Under- and overtreatment might thus occur, even if the expert invests in information and reports it truthfully, as his signal is not always correct.

Second, we allow for sequential treatments. Usually, the game ends when the consumer experiences undertreatment. In reality, however, the interaction between the consumer and the expert is unlikely to end if the consumer’s problem is not solved. A patient, for instance, whose health problem persists after a non-invasive intervention (say acupuncture) might well revisit the expert and ask for a surgical procedure. We assume that after the low-cost treatment failed, the principal’s problem may be solved by applying the high-cost treatment in a second period. Undertreatment causes therefore two types of costs. On one hand, the first period low–cost treatment is wasteful, and, on the other, it delays the solution of the principal’s problem. We measure this delay cost by $1 - \delta$, where δ is the parties’ common discount factor. When the discount rate is large (and δ low), undertreatment involves

certifying physicians specializing in internal medicine) has started the initiative Choosing Wisely that aims to reduce overuse of tests and procedures, see www.choosingwisely.org, accessed on 23/07/2014.

³Kale et al. (2011) look at that the top 5 overused clinical activities across 3 primary care specialties (pediatrics, internal medicine, and family medicine) and report that these activities are common in primary care although they provide little benefit to patients. They conclude that the associated costs of these activities exceed \$ 5 billion.

substantial delay cost, and in the extreme we recover the standard setting considered in the survey by Dulleck and Kerschbamer (2006). Decreasing the discount rate, the other extreme is eventually reached, in which the discount rate is zero (and $\delta = 1$) so that undertreatment involves no delay cost at all.⁴

Third, we connect the credence good problem to the literature on contracting with subjective evaluation. As already mentioned, we assume that treatment choices are verifiable and contractible. This assumption is likely to hold when the consumer needs no specific expertise to identify treatments, as in the aforementioned examples of acupuncture and surgical procedures (when informed consent must be given).⁵ In addition, we suppose that - even though the consumer might be able to observe the outcome of treatment as private information - it is impossible to verify treatment success in court.⁶ As already mentioned, however, the interaction between the consumer and the expert is unlikely to end if the consumer experiences undertreatment. We therefore assume that undertreatment can have consequences for the agent. More precisely, we assume that payments for the low-cost treatment can be made contingent on the principal's subjective report of the outcome, that is, on whether he reports success or failure. Since this report might potentially be misrepresented, a *problem of subjective evaluation* arises.

We analyse whether the principal can choose the payments for treatments in such a way that delegating the choice of treatment to the agent avoids both the *moral hazard problem* and the *problem of subjective evaluation*.⁷ The motivation for this approach is the following. In our model treatments are contractible. Following the Revelation Principle, the contracting problem can then be stated as a direct mechanism design problem. Con-

⁴Liu and Ma (2013) provide a model of medical treatment decisions with treatments potentially applied in sequence. To the best of our knowledge only Dulleck and Kerschbamer (2009) consider a model of credence goods and allow for sequential treatments. Both papers consider only the extreme case in which the discount rate is zero. This precludes delay costs, which - as we will see - play an important role in our setting.

⁵Dulleck and Kerschbamer (2006, p. 16) write that "The verifiability assumption is likely to be satisfied in important credence goods markets, including dental services, automobile and equipment repair, and pest control. For more sophisticated repairs, where the customer is usually not physically present during the treatment, verifiability is often secured indirectly through the provision of ex post evidence. In the automobile repair market, for instance, it is quite common that broken parts are handed over to the customer to substantiate the claim that replacement, and not only repair, has been performed. Similarly, in the historic car restoration market the type of treatment is usually documented step by step in pictures." See also Dulleck and Kerschbamer's discussion on p. 31.

⁶This difficulty has been recognized in the literature. Dulleck and Kerschbamer (2006, p. 32) write that "... 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). In such a situation, a patient may misreport treatment success ...".

⁷There is evidence that financial incentives affect professional recommendations, including studies on health services (e.g. Donaldson and Gerard (1989), Henning-Schmidt et al. (2011), Clemens and Gottlieb (2014)) and financial services (e.g. Mullainathan et al. (2012), Brown and Minor (2013)).

sequently, it is optimal to specify contract terms depending on the expert's reports, subject to the requirement that it must be incentive compatible for him to reveal his information truthfully. Therefore, it is also optimal to delegate the choice of treatment to the expert. A second feature of our approach is similar to Liu and Ma (2013) in that contractual payments are specified over the entire diagnosis and treatment relation. This specification allows the payment for the high-cost treatment in the second period (after failure of the low-cost treatment) to differ from the payment for using high-cost treatment directly in the first period.

Our main result highlights the importance of delay costs. If the discount rate is zero (and $\delta = 1$) so that undertreatment involves no delay cost, then the first-best is *always* implementable by a contract. If, however, the discount rate is positive (and $\delta < 1$), then the first-best is still implementable when effort costs are either high or low. For situations with intermediate effort costs in combination with a relatively uninformative prior, however, only a second-best outcome is implementable. The intuition for this result is based on the interplay of the problems of *moral hazard* and *subjective evaluation*. On the one hand, avoiding the former and endowing the expert with efficient information acquisition incentives requires that his reward in the event of failure of the low-cost treatment should be small enough in order to reflect delay cost. On the other hand, if payments after failure are too low, then the principal has no incentive to reveal the success of the low-cost treatment truthfully and the latter problem cannot be avoided. Consequently, the conflict between the requirements grows with delay costs and the requirements can only be reconciled when these costs vanish.

We turn then to an extension of our model in which it is possible to separate diagnosis and treatment. This is a reasonable assumption in situations in which diagnosis and treatment are to a large extent two independent procedures. For example, prescribing and vending of drugs are independent activities and can therefore be carried out by different agents (doctors and pharmacies). In other situations, however, such a separation is not reasonable, as when treatment is more or less a by-product of diagnosis (like some surgical procedures).⁸ We show that when separation is possible, then the first-best is always im-

⁸Dulleck and Kerschbamer (2006) provide further examples for situations in which separation is likely to be feasible and when it is not. Emons (1997) relates the feasibility of separation to the existence of economies of scope between diagnosis and treatment, while Pesendorfer and Wolinsky (2003) propose a formalization that is based on whether a recommendation identifies uniquely the service to be performed. Separation has already been discussed by Darby and Karni (1973) as a solution to the problem of fraud by experts, because an expert has only an incentive to recommend unnecessary treatment if he also provides it. As observed by Wolinsky (1993, p. 387), "Such arrangements, however, would raise new problems regarding the proper incentives for the diagnostician and hence might not be easily sustainable." In our setting the expert must not only be provided with incentives to give correct advice, but also to exert costly diagnostic effort in the first place.

plementable. The intuition is that separate payments for diagnosis and treatment allow for more flexibility to cope with the joint problems of *moral hazard* and *subjective evaluation*. The trade-off between both problems can be solved in the following way. In the event of failure, the agent making the diagnosis can receive a small payment so that information acquisition is induced, while the agent providing the treatment receives a large amount so that total payments to both agents are high enough to give the principal incentives for truthful reporting.

Although our model is very stylized, it allows us to draw policy implications. In situations in which diagnosis and treatment are jointly provided, it recommends contracts that resemble a *prospective payment system* and contributes hence to the discussion whether a retrospective or a prospective payment system is optimal.⁹ Under a prospective system payments have no link to the real costs of treatments, while under a retrospective system reimbursement is based on real treatment costs. The standard argument in favor of a prospective payment system is that a fixed budget contains health care costs, because it breaks the link between the agent's income and the number of treatments provided. Our argument focusses on a moral hazard problem in providing diagnosis and is therefore different. Our analysis also recommends that contractual payments should be based on treatment protocols, rather than on each treatment in isolation. Payments for the low-cost treatment should include a reimbursement for a potentially necessary high-cost treatment. This mimics the so-called *payments per case* to both individual providers, like general practitioners, or institutional providers, like hospitals.¹⁰ In addition, our model recommends equal markup payments on treatment protocols that are based on the expected costs of a sequence of treatments. Prospective reimbursement systems with this feature are called *cost-neutral*. Cost-neutrality is usually justified by a concern to avoid that patients receive the least costly treatment option. Our argument is again different, since it recognizes the existence of diagnosis costs and is based on a concern to provide incentives for information acquisition. Taking all together, our model provides support for a prospective reimbursement payment system in which reimbursement is per case and cost-neutral. This stands in stark contrast to the standard credence good model which recommends a retrospective payment system with equal markup payments for real treatment costs.

In some situations it is possible to separate diagnosis and treatment, because they are two largely independent procedures. In these cases our analysis provides novel support

⁹Much of the following discussion is based on the definitions provided in Jegers et al. (2002).

¹⁰An important example for payments per case to institutional providers are Diagnostic Related Groups Systems (DRG-system) for hospitals. This system has been used in the U.S. (Medicare), Canada, or Germany, among others (see Jegers et al. (2002)). An example for payments per case to individual providers is a so-called integrated system, like the British general practitioner fundholder system. Here primary care doctors might get a payment in advance and in return have to provide comprehensive health care that might require buying hospital care or drugs when needed.

for separation, complementing existing explanations. For example, in their discussions of separation between doctors and pharmacies. Darby and Karni (1973) and Dulleck and Kerschbamer (2006, p. 8) focus on the moral hazard problem and argue that breaking up the joint provision of diagnosis and treatment avoids overtreatment. In contrast, in our framework separation alleviates not only the moral hazard problem but also the tension between this problem and the problem of subjective evaluation.

Indeed, there is reason to believe that separation can contribute to avoid conflicts of interest that may lead to over-utilization of treatments and result in large health care costs. In the U.S. the so-called Stark Law governs physician self-referral for Medicare and Medicaid patients, where self-referral is defined as “the practice of a physician referring a patient to a medical facility in which he has a financial interest, be it ownership, investment, or a structured compensation arrangement.”¹¹ Intensity-modulated radiation therapy (IMRT), a treatment for prostate cancer, is an exception from the general Medicare prohibition of self-referral. A recent report to the U.S. Congress explains that several other less costly and equally appropriate alternative treatments to IMRT are available. The report looked at providers who began to self-refer to IMRT during the study. These providers increased the percentage of prostate cancer patients referred for IMRT from 37% to 56% after they began to self-refer. Moreover, during the period studied the number of IMRT services performed by self-referring providers and its associated costs increased.¹²

Related Literature

The literature on credence goods focusses on the avoidance of fraud. Since the result of diagnosis is unobserved, opportunistic advice might result in inappropriate treatment. We follow a strand of the literature that assumes treatment decisions are observable. As mentioned earlier, fraud might then occur when treatments have different markups, as the expert might misrepresent the required treatment.¹³ A second strand of literature assumes that the treatment provided is not observable. This gives rise to a different type of fraud, in which the low-cost treatment is provided but misrepresented as a high-cost one.¹⁴ In addition, a social loss from the problem remaining insufficiently treated may arise.¹⁵

These models assume that the expert is perfectly informed, without incurring any costs. Consequently, they abstract from the moral hazard problem we consider.¹⁶ When such a

¹¹See www.starklaw.org, accessed on 23/07/2014.

¹²United States Government Accountability Office (2013).

¹³See Darby and Karni (1973), Emons (1997, 2001), De Jaegher and Jegers (2001), Frankel and Schwarz (2014), Fong et al. (2014).

¹⁴See Darby and Karni (1973), Pitchik and Schotter (1987, 1993), Wolinsky (1993, 1995), Sülzle and Wambach (2005), Alger and Salanié (2006), Hyndman and Ozerturk (2011).

¹⁵See Fong (2005) and Liu (2011).

¹⁶Dulleck and Kerschbamer (2006) review this literature in a unified model. Our model generalizes their

moral hazard problem is considered, equal markups no longer achieve the first–best and the question how the credence good environment should be designed arises. The theoretical literature considering this question is very small.¹⁷

Similar to us, Pesendorfer and Wolinsky (2003) assume that success and failure of treatment are unobservable and unverifiable, but they do not consider the problem of subjective evaluation. In addition, the problem of overtreatment cannot appear, treatments are equally costly, and experts have no incentive to misrepresent the required treatment, once it is diagnosed. This isolates the problem to provide incentives to invest in costly diagnostic effort, when contrary to us the diagnostic signal is assumed to be perfectly precise. Pesendorfer and Wolinsky explore how competition between experts can allow consumers to obtain multiple diagnoses and whether this competition can provide experts with effort incentives. Their main result is that the equilibrium is inefficient. In contrast, we provide conditions under which the first–best can be reached. In an extension, Pesendorfer and Wolinsky (2003) also discuss briefly the case when it is possible to separate diagnosis and treatment. Interestingly, they show that the outcome remains inefficient, while our extension to such a setting shows that the first–best is always implementable. In this sense, our result emphasises the potential of multiple agents to *specialize* in diagnosis and treatment rather than to *compete* with one another, as in Pesendorfer and Wolinsky.¹⁸

Dulleck and Kerschbamer (2009) investigate the incentives of experts to exert diagnostic effort and to report truthfully when they face competition from chain stores or discounters. Their analysis has thus a very different focus from ours. In their model, however, effort is costly and treatments can be applied in sequence. The main modeling differences are that in their setting the expert’s signal is perfectly precise, that there are no delay costs and, most importantly, that success and failure of treatment are observable and verifiable. A byproduct of their analysis is to describe contracts that provide effort incentives. This is consistent with our result that the first–best can be achieved when there are no delay costs. Our analysis goes beyond their result and shows that when failure is not verifiable

setting in the three dimensions explained above. The limiting case of our model when (i) the expert’s signal is costless as well as perfectly precise; (ii) treatments cannot be applied in sequence (that is, the discount factor δ vanishes); and (iii) contractual payments are restricted to depend only on the choice of treatment, recovers the setting in Dulleck and Kerschbamer (2006) in which their assumptions H (Homogeneity), C (Commitment), and V (Verifiability) hold and their Lemma 1 applies.

¹⁷Bonroy et al. (2013) also analyse a credence good problem and look at incentives for costly diagnostic effort. The problem considered is different from ours, however, because effort is observable. The papers by Demski and Sappington (1987), Taylor (1995), and Malcomson (2004) investigate costly diagnostic effort in settings different from the credence goods environment with observable treatments that we consider.

¹⁸Notice that this specialization result is different from (vertical) specialization results in the second strand of literature mentioned above, in which treatment decisions are unobservable. In Wolinsky (1993) or Alger and Salanié (2006) some experts provide only the minor treatment, while others specialize in the major treatment.

and delay costs are introduced, a tension between the problems of *moral hazard* and *subjective evaluation* arises, so that the first–best can no longer be achieved for some parameter combinations.

Finally, our paper is related to the literature on subjective evaluation, which addresses the problem of providing of effort incentives for the agent in situations where only the principal privately observes some performance measure. This applies to our context because the expert has to invest effort in identifying the appropriate treatment, and only the consumer learns the outcome of the treatment. Some part of the literature studies subjective evaluations in models of repeated interactions, where intertemporal incentives play a key role.¹⁹ Our model is more closely related to MacLeod (2003), who like us considers a single interaction between a principal and an agent. He shows that effort incentives can be created only if the contract specifies some ex post inefficiencies that are formally equivalent to ‘money burning’, i.e. payments to a passive third party. In contrast, our analysis is confined to budget balanced contracts and excludes third-party payments. Nonetheless, at least for some range of parameter combinations, the agent can be incentivized to exert diagnosis effort. The reason is that after failure of the first treatment, he can be contractually obliged to perform a costly second treatment. This creates a possibility to punish the agent for failure and so he can be induced to acquire information about the appropriate first period treatment.

This paper is organized as follows. The next section describes the credence goods problem and our assumptions on observability and contracts. In Section 3 we characterize the first–best decision. Section 4 analyses optimal contracts and establishes our main result. In Section 5 we extend our model to a situation in which it is possible to separate diagnosis and treatment. Finally, we conclude in Section 6.

2 The Model

The Credence Good Problem

Consider the by now standard credence good problem. An expert potentially knows more about the quality of a good or service that a consumer needs than the uninformed consumer himself.²⁰ We also refer to the consumer as the principal and to the expert as the agent.

More precisely, the consumer needs one of two services (or treatments), depending on his type. If the principal has the minor problem θ_L , then the low-cost treatment T_L is sufficient. The problem might, however, be important, denoted by θ_H . In this case the

¹⁹See e.g. Baker, Gibbons and Murphy (1994), Fuchs (2007), Levin (2003), Pearce and Stacchetti (1998).

²⁰See the survey by Dulleck and Kerschbamer (2006).

high-cost treatment T_H is needed, as the low-cost service does not help. This formalizes one type of choice in health care that is often considered to lead to wasteful expenditures. Discussing waste in health care, Fuchs (2009) gives the following example for a choice between a high-cost procedure and a less expensive alternative: “high-cost drug-eluting stents may be the better choice for some patients, but others would do just as well with less expensive bare-metal stents” (p. 2481). Formally, the principal’s gross utility u^P depends on his type $\theta \in \{\theta_L, \theta_H\}$ and the treatment $T \in \{T_L, T_H\}$ according to²¹

$$u^P(\theta, T) = \begin{cases} 0 & \text{if } \theta = \theta_H \text{ and } T = T_L \\ v > 0 & \text{otherwise.} \end{cases}, \quad (1)$$

We refer to the combination (θ_H, T_L) as *undertreatment*: the consumer discovers ex post that he has the major problem and that the service T_L is insufficient. When the minor problem is solved through the high-cost treatment in the combination (θ_L, T_H) , we speak of *overtreatment*, as the low-cost service would have been sufficient.

The principal is uncertain as to which service is the correct one. But he knows the prior probability

$$\text{Prob}(\theta_L) = 1 - \text{Prob}(\theta_H) = q. \quad (2)$$

The outside option of the principal is not to be treated at all, giving zero utility to both principal and agent.

Also the expert a priori only knows (2). But he can acquire additional information about the principal’s problem to identify the appropriate treatment. To do so he needs to exert effort at a cost $c \geq 0$, which enables him to privately observe a signal $s \in \{s_L, s_H\}$ about the principal’s problem θ . The signal is correct with probability σ , i.e.

$$\text{Prob}(s_L|\theta_L) = \text{Prob}(s_H|\theta_H) = \sigma, \quad \text{Prob}(s_L|\theta_H) = \text{Prob}(s_H|\theta_L) = 1 - \sigma, \quad (3)$$

with $\sigma > 1/2$. If the agent exerts no effort, he incurs no cost but learns nothing. These assumptions generalize the previous literature on credence goods which assumed either $c = 0$ or $\sigma = 1$ or both. The effort cost c can also be interpreted as the opportunity costs of time. Physicians often complain that changes in reimbursement oblige them to see more patients per day, thereby making it more difficult to conduct proper diagnosis.²²

Once the principal was treated in the first period and experienced undertreatment, in a second period he can request the high-cost treatment, which will then solve his problem.

²¹The assumption that when the problem is important the low-cost treatment never solves it, while the high-cost treatment solves it with probability one, is standard in the credence goods literature. It seems, however, stronger than needed. Suppose the low-cost treatment has a small and the high-cost treatment a high probability of solving the major problem respectively. Then it can be shown that our main result that the first-best can be reached, provided delay costs are absent, is still true.

²²See e.g. Sandeep Jauhar, “Busy Doctors, Wasteful Spending,” The New York Times 20, July 2014.

We introduce the discount factor $\delta \in [0, 1]$, in order to capture that the principal prefers his problem to be solved in period 1 rather than in period 2. A special case is when $\delta = 0$ and only one treatment can be applied. This might be because the treatment is urgent (the principal is extremely impatient) or the first treatment is irreversible. As explained in the literature review, the literature on credence goods focusses almost exclusively on the case $\delta = 0$ and is thus generalized through our model. For simplicity we assume that the principal and the agent share a common discount factor.

Abusing notation we indicate also the treatment costs by T_L and T_H . Unless explicitly stated otherwise, we assume that $v > T_H > T_L \geq 0$, $q \in (0, 1)$, $c > 0$, $\delta > 0$, and $\sigma > 1/2$. The values of these parameters are common knowledge.

Observability and Contracts

We assume that the expert's choice of treatment is observable. As mentioned in the Introduction, this is likely to hold when no specific expertise is required to identify treatments. Therefore, a contract can stipulate that he selects treatment T_H for a payment p_H in the first period. Since it is commonly known that this solves the principal's problem, under this contract the principal–agent relation ends at the end of the first period.

Also, a contract can specify that the agent selects treatment T_L in the first period. We assume, however, that success and failure are not publicly observable. The principal privately learns whether this treatment has been successful or not at the end of the first period. Therefore, the payment for treatment T_L can only depend on a *subjective evaluation* $R \in \{S, F\}$ of the principal, where S indicates "success" and F "failure". Thus, if the agent selects treatment T_L , he receives the payment p_{LS} if the principal reports S and p_{LF} otherwise. This captures for example health problems for which it is difficult to measure treatment success objectively. Dulleck and Kerschbamer (2006) mention the difficulty of proving the absence of pain as an example. It is less appropriate when an objective test exists, like in the case of cancer screening.

Since $v > T_H$, a contract optimally entitles the principal to treatment T_H in the second period upon failure of treatment T_L in the first period. Without loss of generality, there are no additional payments for the second treatment. This means that the payment p_{LF} is the reimbursement for treatment T_L in the first period and T_H in the second period.²³ As the principal can misreport a successful treatment, for completeness we also have to specify his utility for the case where he demands T_H even though his problem has been solved by T_L in the first period. We assume that in this case a second treatment in period 2 does not affect

²³If the principal has to pay p'_{LF} for the first and p'_2 for the second treatment, this is equivalent to a single payment $p_{LF} = p'_{LF} + \delta p'_2$ after the first treatment.

the principal's gross utility, i.e. his gross payoff remains v at the end of period 1.²⁴

Finally, the principal may wish the agent to exert diagnosis effort before a treatment is selected. Yet, in addition to the problem that success and failure of the low-cost treatment are not publicly observable, this creates a *moral hazard problem*: neither the public nor the principal observe whether the agent invests effort in information acquisition, and if so which signal he observes. Therefore, if the principal prefers a costly diagnosis, he has to delegate the choice of treatment to the expert and to choose the payments

$$p \equiv (p_H, p_{LS}, p_{LF}) \quad (4)$$

in such a way that they provide the incentive to acquire information about θ . Indeed, under an optimal contract, the Revelation Principle requires the agent to report the observed signal truthfully and the principal to commit himself to a treatment strategy contingent on the agent's report. In line with the Delegation Principle (see, e.g., Holmström (1984) and Alonso and Matouschek (2008)) this is equivalent to a contract that delegates the choice of treatment to the agent.²⁵

The contracting relation proceeds in the following stages:

1. Nature determines the principal's type $\theta \in \{\theta_L, \theta_H\}$. Neither the principal nor the agent observes the realization of θ . They both know only the a priori probabilities as given by (2).
2. The principal signs a contract with the agent. This specifies some payments p . In addition, the principal can either delegate the choice of treatment to the agent or he can demand some first-period treatment $T \in \{T_L, T_H\}$. Our assumptions ensure that a positive net surplus can be achieved by an appropriate contract, on which both parties will agree.
3. If the treatment is already specified in the contract, the agent selects the mandated treatment. Otherwise, if the choice of treatment is delegated, he decides about whether or not to invest effort in information acquisition. After investing he privately observes a signal s , updates his prior beliefs according to (3), and then chooses some first-period treatment T . The agent's effort decision is not observable. Without information acquisition the agent directly selects T without observing a signal.

²⁴Constant gross utility is the conservative assumption to make, as if the principal's gross utility were to decline, incentives for misreporting were reduced. On the other hand, increasing gross utility is not in line with the basic assumptions of the credence goods problem, in which the high-cost treatment does not yield higher gross utility than the low-cost treatment given that both solve the problem.

²⁵Delegating the choice of treatment assumes that after hearing the diagnosis the principal is committed to undergo the recommended treatment. Without such a commitment the principal could also leave without payment after hearing the diagnosis and request the recommended treatment from another agent.

	θ_L	θ_H
T_L	$v - T_L$	$-T_L + \delta(v - T_H)$
T_H	$v - T_H$	$v - T_H$

Table 1: Surplus net of treatment cost

4. After a first-period treatment T_H the contracting relation ends and the principal pays p_H . If treatment T_L has been selected, the principal privately observes whether his problem has been solved. If he reports “success” he pays p_{LS} and the relation ends; if he reports “failure” he pays p_{LF} and receives treatment T_H in the second period.

Thus, the principal can either keep authority over the selection of treatment or he can delegate the treatment decision to the agent. If he is confident that he can identify the appropriate treatment, he can follow his judgement. Otherwise he has the opportunity of letting the expert determine the treatment decision. This seems to be in line with patient preferences. When the choice of the appropriate treatment is uncertain patients prefer to delegate the final decision to their physician, rather than making the decision themselves.²⁶

3 First-Best Treatment Strategies

Before analyzing the optimal contract between the principal and the agent, we consider the first-best outcome. Suppose the principal is able to acquire information and to perform the appropriate treatment with the same cost as the expert. Since the principal himself effectively takes over the role of the expert, both the problem of subjective evaluation and the moral hazard problem of investing effort then disappear. The principal maximizes the overall surplus and the result is the first-best decision.

As explained before, when treatment T_L fails in period 1, it is optimal to choose T_H in the second period. Therefore, the different combinations of treatments T and types θ imply the surplus net of treatment cost given in Table 1.

At the beginning of the first period the principal can either choose the treatment based on prior information only or he can invest in information before making his choice. Thus there are three possible treatment strategies. First, the principal can choose treatment T_H without acquiring information. This yields the net surplus

$$S_H^* \equiv v - T_H, \tag{5}$$

²⁶See Pauline W. Chen, M.D., “Letting Doctors Make the Tough Decisions,” The New York Times 11, August 2011.

because the treatment is always successful. Yet, with probability q it involves overtreatment. Second, also based on prior information only, the principal can first try the low-cost treatment and correct this choice later when needed. We call this the *trial-&-error* strategy. It yields the net surplus

$$S_{T\&E}^* \equiv q(v - T_L) + (1 - q) [-T_L + \delta(v - T_H)], \quad (6)$$

because the problem is solved with the a priori probability q , whereas with probability $1 - q$ it turns out that T_L results in undertreatment so that the high-cost treatment becomes necessary in period 2.

Finally, if the principal exerts effort in diagnosis and chooses treatment T_i upon observing signal s_i , he obtains²⁷

$$S_i^* \equiv q [\sigma(v - T_L) + (1 - \sigma)(v - T_H)] \\ + (1 - q) [\sigma(v - T_H) + (1 - \sigma)(-T_L + \delta(v - T_H))] - c. \quad (7)$$

Indeed, with probability q the problem is minor and overtreatment occurs only if the signal is incorrect. With probability $1 - q$ the problem is major, and when the signal is incorrect the treatment decision must be corrected later. These expected benefits are reduced by the information cost c . More precise information is beneficial because

$$\partial S_i^* / \partial \sigma = (1 - \delta)(1 - q)(v - T_H) + q(T_H - T_L) + (1 - q)T_L > 0. \quad (8)$$

In fact, when the signal becomes perfectly precise as $\sigma \rightarrow 1$, the information acquisition strategy reduces the likelihood of both overtreatment and undertreatment to zero.

From the payoffs in (5)–(7) we can now derive the first-best treatment strategy. If treatment choice is based on prior information only, the trial-&-error strategy is at least as good as choosing the high-cost treatment if $S_{T\&E}^* \geq S_H^*$, which is equivalent to

$$q \geq q^* \equiv \frac{(1 - \delta)(v - T_H) + T_L}{(1 - \delta)(v - T_H) + T_H}. \quad (9)$$

Clearly, our assumptions imply that $q^* \in (0, 1)$. Intuitively, the trial-&-error strategy is the more attractive, the more the principal is concerned with overtreatment and the less he cares about undertreatment.

Investing in information is optimal if it is at least as good as choosing any of the two strategies based on prior information, i.e. if $S_i^* \geq S_H^*$ and $S_i^* \geq S_{T\&E}^*$. These two conditions

²⁷It is easy to see that investing in diagnosis and then ignoring the information cannot be optimal. Therefore, selecting T_i after signal s_i is optimal whenever information acquisition yields a higher expected surplus than a decision based on prior information. The formal argument is analogous to the proof of Lemma 3 below and is omitted here.

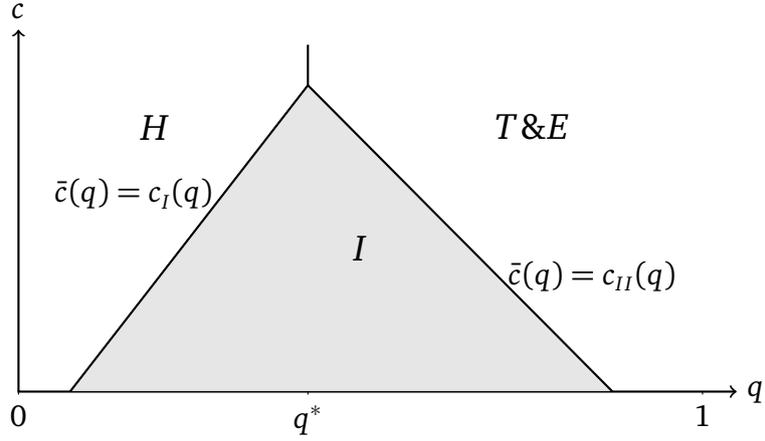


Figure 1: *First-best treatment strategies*

are satisfied if and only if

$$c \leq c_I(q) \equiv (T_H - T_L) [q(2\sigma - 1) + 1 - \sigma] - (1 - \sigma)(1 - q) [\nu(1 - \delta) + \delta T_H] \quad (10)$$

and

$$c \leq c_{II}(q) \equiv (T_H - T_L) [q(2\sigma - 1) - \sigma] + \sigma(1 - q) [\nu(1 - \delta) + \delta T_H]. \quad (11)$$

Therefore, as long as

$$c \leq \bar{c}(q) \equiv \min[c_I(q), c_{II}(q)], \quad (12)$$

the principal optimally invests in information acquisition before taking a treatment decision.

It is easily verified, that $c_I(\cdot)$ and $c_{II}(\cdot)$ are linear in q with $\partial c_I(q)/\partial q > 0$ and $\partial c_{II}(q)/\partial q < 0$. Further, for q^* as defined in (9) we have

$$\bar{c}(q^*) = c_I(q^*) = c_{II}(q^*) = \frac{(2\sigma - 1)(T_H - T_L) [(1 - \delta)(\nu - T_H) + T_L]}{\nu(1 - \delta) + \delta T_H} > 0. \quad (13)$$

As illustrated in Figure 1, this implies that the critical level of information costs $\bar{c}(\cdot)$ is linearly increasing in q for $q < q^*$ and decreasing for $q > q^*$ so that $\bar{c}(\cdot)$ is maximized by q^* . Moreover, as $\bar{c}(q^*) > 0$, by (12) information acquisition is the optimal strategy for some interval $Q(c)$ of q -values with q^* in its interior whenever $c \leq \bar{c}(q^*)$. As c decreases, $Q(c)$ expands so that information acquisition becomes attractive for a larger range of parameter combinations.

The following proposition summarizes the first-best treatment strategy:

Proposition 1 *The first–best solution has the following properties:*

- (a) *If $c \geq \bar{c}(q)$ and $q \leq q^*$, it is optimal to choose the high-cost treatment without diagnosis.*
- (b) *If $c \geq \bar{c}(q)$ and $q \geq q^*$, it is optimal to choose the trial-&-error strategy, i.e. the low-cost treatment without diagnosis, followed by the high-cost treatment in case of failure.*
- (c) *If $c \leq \bar{c}(q)$, it is optimal to exert effort in diagnosis and choose the treatment contingent on the information revealed.*

Figure 1 illustrates Proposition 1. Information acquisition constitutes the first–best strategy in the gray shaded area. Outside this area, the first–best outcome requires the high–cost treatment if $q \leq q^*$ and the trial-&-error strategy otherwise. As can be seen from the shape of the gray shaded area, the more diffuse the prior, the less is known about the success of treatments and the higher the incentives for information acquisition. As the prior becomes more precise, the risk of overtreatment with the high-cost treatment or of undertreatment with the trial-&-error strategy declines, because $\partial c_I(q)/\partial q > 0$ and $\partial c_{II}(q)/\partial q < 0$ respectively. Indeed, when the prior becomes perfectly precise as $q \rightarrow 0$ or $q \rightarrow 1$, the treatment decisions based on prior information entail no risk of over- or undertreatment and information acquisition is not needed to take the correct treatment decision. Moreover, the overall incentives to invest in information (measured e.g. by the altitude of the gray shaded triangle) are the higher, the more important it is to avoid over- and undertreatment and the more precise the signal is that can be acquired.

Before analysing optimal contracts under the informational assumptions of Section 2, it may be useful to point out that despite the non–observability of diagnosis effort the first–best treatment strategy can easily be implemented by a contract as long as success and failure of a treatment are publicly observable. Payments then can be made directly contingent on the treatment outcome: If already the first–period treatment is successful, the expert receives the payment $p = v - k$ in period 1; otherwise the first period payment is reduced to $p = \delta v - k$ and the expert is contractually obliged to administer the high-cost treatment in period 2. The parameter k is some constant that determines the division of expected surplus between both parties.

With these payments, the agent’s expected profit net of expected treatment costs is $S_H^* - k$ if he chooses the high cost treatment in period 1. If he adopts the trial-&-error strategy, he gets $S_{T\&E}^* - k$; and if he invests in information before selecting the first–period treatment he gets $S_I^* - k$. Thus, by the above payments the agent becomes the residual claimant and he will choose the treatment strategy that maximizes the first–best surplus.

4 Optimal Contracts

We now study the optimal contract between the consumer and the expert. Since both parties are risk-neutral, they agree at the contracting stage to maximize their joint surplus. To derive the optimal contract, we can therefore focus on contracts that maximize the net surplus. Under our assumptions a positive net surplus can be achieved by an appropriate contract, on which both parties will agree. The actual division of surplus depends on market conditions and can be determined by some upfront payments or by adjusting the payments p in (4) appropriately.²⁸ If, for example, there are several competing experts and the consumer has all the bargaining power at the contracting stage, he can appropriate the entire joint net surplus in this way.

We first investigate the possibility of implementing the first-best outcome through a contract. Trivially, this is possible for all parameter combinations described in part (a) of Proposition 1, where the first-best solution is to choose the high-cost treatment without prior diagnosis. In this situation the agent can simply be contractually obliged to select treatment T_H for a payment p_H . The principal's and the agent's payoffs from such a contract are

$$U_H(p) \equiv v - p_H, \quad V_H(p) \equiv p_H - T_H. \quad (14)$$

The agent's reimbursement can be set equal to $p_H = T_H + k$, where k is some constant that can be adjusted to divide the joint surplus $S_H^* = U_H(p) + V_H(p) = v - T_H$. In the extreme cases, if the principal has all the bargaining power, $k = 0$ and the agent's net payoff is zero; if the agent has all the bargaining power, $k = v - T_H$ and so the principal's net payoff is zero.

Next consider the case where the trial-&-error strategy is optimal in the first-best, i.e. where part (b) of Proposition 1 applies. In this case contracting is slightly complicated by the fact that the consumer privately observes success or failure of the first-period treatment. Optimal contract design requires that he publicly reports his information (see Myerson (1986)). Further, by the Revelation Principle (Myerson (1979)) there is no loss of generality in considering only contracts under which reporting is truthful. The following lemma describes the restrictions that the *truthful reporting* requirement imposes on the payments p :

Lemma 1 *The principal reports success and failure truthfully after treatment T_L if and only if*

$$p_{LF} - \delta v \leq p_{LS} \leq p_{LF}. \quad (15)$$

²⁸As our analysis below shows, incentive effects depend only on payment differences for different treatments. Therefore, the *level* of payments can be adjusted to reflect market power.

Proof: If treatment T_L was successful, the principal's payoff from truthful reporting at the end of period 1 is $v - p_{LS}$. If he reports F , he gets $v - p_{LF}$. Therefore, the second inequality in (15) ensures that he reports truthfully. If the treatment T_L failed, the principal's payoff from reporting F is $-p_{LF} + \delta v$ and his payoff from reporting S is $-p_{LS}$. By the first inequality in (15), he therefore reports truthfully. Q.E.D.

Intuitively, reporting success must be cheaper for the principal than reporting failure. The difference, however, cannot be larger than the gains from receiving the high-cost treatment in the second period. Given that the condition for truthful reporting in (15) holds, at the contracting stage the principal's and the agent's payoffs from the trial-&-error strategy are

$$\begin{aligned} U_{T\&E}(p) &\equiv q(v - p_{LS}) + (1 - q)(\delta v - p_{LF}), \\ V_{T\&E}(p) &\equiv q(p_{LS} - T_L) + (1 - q)(p_{LF} - T_L - \delta T_H), \end{aligned} \tag{16}$$

because with probability q the first-period treatment T_L is successful and with probability $1 - q$ it fails, requiring the high-cost treatment in period 2. For any p satisfying (15), the contracting parties can achieve the first-best joint surplus $S_{T\&E}^* = U_{T\&E}(p) + V_{T\&E}(p)$.

Obviously, it is not problematic to write a contract complying with the incentive-compatibility constraint (15). Since the interval $[p_{LF} - \delta v, p_{LF}]$ is non-empty, it is always possible to choose a price p_{LS} within this interval. In particular, consider a contract with *equal markup* payments as in Dulleck and Kerschbamer (2006) so that the expert's net payoff is independent of treatment costs. Such a contract is defined by the property that

$$p_H - T_H = p_{LS} - T_L = p_{LF} - (T_L + \delta T_H) = k, \tag{17}$$

for some markup k that can be adjusted to determine the division of the joint surplus. As $v > T_H > 0$, it is easy to see that these payments satisfy (15). With equal markups the principal fully bears the cost of an additional treatment, and therefore he always reports truthfully.

The following proposition summarizes our observations so far:

Proposition 2 *Let $c \geq \bar{c}(q)$, i.e. the first-best requires no investment in diagnosis. Then there exists an optimal contract that implements the first-best solution.*

We now turn to the more interesting and also more complicated case where the first-best treatment strategy involves diagnosis effort by the expert, as in part (c) of Proposition 1. Implementing this strategy by a contract requires not only that the principal truthfully reports the outcome of the low-cost treatment, but also that the agent is motivated to invest c in information acquisition before choosing a treatment T . Consider a contract satisfying

(15) so that the first of these two requirements is fulfilled. Then, if the agent exerts effort and chooses treatment T_i upon observing signal s_i , the expected payoffs of the principal and the agent are²⁹

$$U_I(p) \equiv q [\sigma(v - p_{LS}) + (1 - \sigma)(v - p_H)] + (1 - q)[\sigma(v - p_H) + (1 - \sigma)(\delta v - p_{LF})], \quad (18)$$

$$V_I(p) \equiv q [\sigma(p_{LS} - T_L) + (1 - \sigma)(p_H - T_H)] + (1 - q) [\sigma(p_H - T_H) + (1 - \sigma)(p_{LF} - T_L - \delta T_H)] - c, \quad (19)$$

and their joint surplus is $S_I^* = U_I(p) + V_I(p)$.

When the expert receives the authority to select a treatment, he will exert diagnosis effort only if this gives him a higher payoff than selecting a treatment based on prior information only. Thus the contractual payments have to satisfy the *effort incentive* constraint

$$V_I(p) \geq \max [V_H(p), V_{T\&E}(p)]. \quad (20)$$

This constraint is equivalent to requiring that p simultaneously solves

$$c \leq c_A(p, q) \equiv (T_H - T_L - p_H) [q(2\sigma - 1) + 1 - \sigma] - (1 - \sigma)(1 - q) (\delta T_H - p_{LF}) + q\sigma p_{LS} \quad (21)$$

and

$$c \leq c_B(p, q) \equiv (T_H - T_L - p_H) [q(2\sigma - 1) - \sigma] + \sigma(1 - q) (\delta T_H - p_{LF}) - q(1 - \sigma)p_{LS}. \quad (22)$$

Therefore, condition (20) can also be written as

$$c \leq \tilde{c}(p, q) \equiv \min[c_A(p, q), c_B(p, q)]. \quad (23)$$

In what follows we say that a treatment with diagnosis effort is *implementable* by a contract with payments p if p satisfies both the principal's truthful reporting requirement (15) and the agent's effort incentive constraint (20), or equivalently (23). The following result characterizes the parameter combinations under which such a contract is feasible.

Proposition 3 *There exist payments p that implement diagnosis effort by the expert if and only if*

$$c \leq \hat{c}(q) \equiv q(1 - q)(2\sigma - 1)\delta T_H. \quad (24)$$

²⁹It is easy to see that selecting T_i after signal s_i is optimal for the agent whenever (20) holds, because investing $c > 0$ in diagnosis and then ignoring the information cannot be optimal. The formal argument is analogous to the proof of Lemma 3 below and is omitted here.

If this condition holds, then diagnosis effort is implementable in particular by the payments

$$\hat{p}_H = T_H + k, \quad \hat{p}_{LS} = \hat{p}_{LF} = T_L + (1 - q)\delta T_H + k, \quad (25)$$

for some constant k , with $0 \leq k \leq S_I^*$.

Proof: Let \hat{p} maximize $\tilde{c}(p, q)$, as defined in (23), subject to the truthful reporting condition (15). Then obviously, diagnosis effort is implementable if and only if $c \leq \tilde{c}(\hat{p}, q) = \tilde{c}(q)$.

If $c_A(p, q) > c_B(p, q)$ then $\tilde{c}(p, q) = c_B(p, q)$ is increasing in p_H by (22) because $q(2\sigma - 1) - \sigma < 0$. If $c_A(p, q) < c_B(p, q)$ then $\tilde{c}(p, q) = c_A(p, q)$ is decreasing in p_H by (21) because $q(2\sigma - 1) + 1 - \sigma > 0$. Therefore, $\tilde{c}(p, q)$ is maximized by p_H if $c_A(p_H, p_{LS}, p_{LF}, q) = c_B(p_H, p_{LS}, p_{LF}, q)$. This yields

$$\hat{p}_H = T_H [1 - \delta(1 - q)] - T_L + qp_{LS} + (1 - q)p_{LF}. \quad (26)$$

Since

$$\tilde{c}(\hat{p}_H, p_{LS}, p_{LF}, q) = q(1 - q)(2\sigma - 1)(\delta T_H + p_{LS} - p_{LF}), \quad (27)$$

is increasing in p_{LS} and decreasing in p_{LF} , it is maximized subject to (15) by setting $\hat{p}_{LF} = \hat{p}_{LS}$. As $\tilde{c}(\hat{p}_H, \hat{p}_{LS}, \hat{p}_{LF}, q) = q(1 - q)(2\sigma - 1)\delta T_H$, this proves (24). Finally, we obtain from (26) for $\hat{p}_{LF} = \hat{p}_{LS}$ that

$$\hat{p}_H - \hat{p}_{LS} = \hat{p}_H - \hat{p}_{LF} = T_H [1 - \delta(1 - q)] - T_L, \quad (28)$$

which is equivalent to (25).

Q.E.D.

Since the principal reports truthfully and because T_L additionally requires T_H with probability $1 - q$ in the next period, the prices \hat{p} in (25) can be interpreted as equal markups on the *expected* treatment costs of choosing T_H or T_L before receiving information. This has two implications. First, these payments make the principal indifferent between reporting success and failure after a successful low-cost treatment inducing hence truthful reporting, as (15) holds. Second, they have the property that they equalize the agent's payoffs from treatment choices based on prior information alone, that is $V_{T\&E}(\hat{p}) = V_H(\hat{p})$. Thus, once the principal prefers information acquisition to one of the treatment choices based on prior information, he also prefers it to the other and this happens whenever information acquisition is cheap enough so that (24) holds. For a given combination of q and c the prices \hat{p} are not necessarily the only ones that implement effort. But they are chosen so that even for the highest cost $c = \tilde{c}(q)$ effort is implemented. In other words, they maximize the range of parameter combinations for which effort is induced.

For $\delta = 0$ the payments in (25) become the equal markup payments in (17). But in this special case $\tilde{c}(q) = 0$ and hence effort is only implemented when it is costless. This raises the important question whether equal markup payments can implement effort for $\delta > 0$. As

we have seen above, the payments in (17) give the principal no incentive to misreport the outcome of the low-cost treatment. Yet, as the following result shows, they fail to provide incentives for the expert to invest in costly information.

Proposition 4 *Equal markup payments, as defined in (17), implement a treatment with diagnosis effort if and only if information acquisition is costless, i.e. $c = 0$.*

Proof: Since equal markup prices satisfy the truthful reporting constraint, it remains to check whether they satisfy the agent's effort constraint. Inserting the prices in (17) into (21) and (22) yields $c_A(p, q) = c_B(p, q) = 0$. Therefore, they satisfy (23) if and only if $c = 0$. Q.E.D.

The intuition for why equal markups do not induce information acquisition is closely related to their virtue in the standard model (Dulleck and Kerschbamer (2006)). In that setting the diagnostic effort of the agent is costless and the only issue is to provide him with the appropriate incentives to reveal his information. If markups are unequal and higher, say, for the high-cost treatment, then the expert has an incentive to always recommend this treatment, even if the low-cost treatment would have been sufficient. So the expert has to be indifferent. But if he is indifferent gaining some markup k with each treatment, then, by construction, both treatment choices based on prior information alone yield k and information acquisition does not pay, because with such a strategy he also obtains k but has to pay the cost of information c .

Under our assumptions we have $\tilde{c}(q) > 0$ for all $q \in (0, 1)$. Therefore, for any $q \in (0, 1)$ diagnosis effort is implementable if c is sufficiently small. Diagnostic effort is easier to implement (or equivalently \tilde{c} is the higher), the higher the precision of information σ , the more diffuse the prior and the larger the scope for combined treatments as measured by δ . While the first two are roughly in line with the first-best outcome, the latter plays an important role. On one hand, we already mentioned that the extreme case of $\delta = 0$ implies $\tilde{c}(q) = 0$ so that effort is only implementable when it is costless. On the other hand, as we will see next, the other extreme of $\delta = 1$ yields an efficiency result.

Proposition 5 *If and only if $\delta = 1$, diagnosis effort can be contractually implemented for all parameter combinations for which it is optimal in the first-best. That is,*

$$\{c, q | c \leq \tilde{c}(q)\} \subseteq \{c, q | c \leq \hat{c}(q)\}$$

if and only if $\delta = 1$.

Proof: From (9), (13) and (24) we obtain that

$$\hat{c}(q^*) - \tilde{c}(q^*) = -\frac{(1 - \delta)(2\sigma - 1)v(T_H - T_L) [(1 - \delta)(v - T_H) + T_L]}{[v(1 - \delta) + \delta T_H]^2}. \quad (29)$$

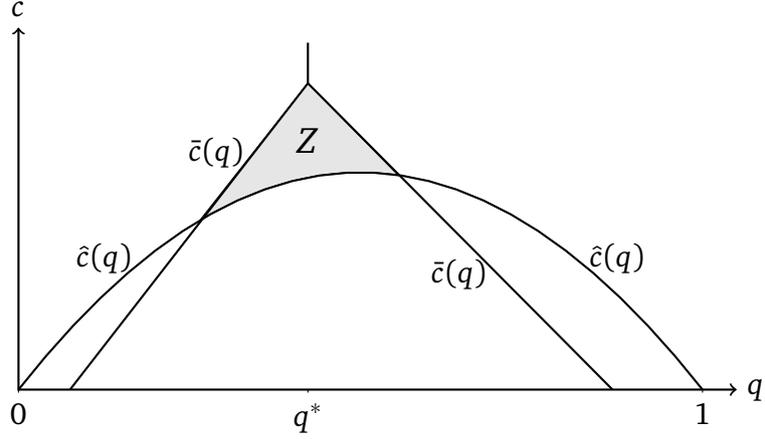


Figure 2: *Non-implementability of first-best*

Further, by (10),(11), and (24) we have

$$\hat{c}(0) - \bar{c}(0) = \hat{c}(0) - c_I(0) = (1 - \sigma) [(1 - \delta)(v - T_H) + T_L] \geq 0, \quad (30)$$

$$\hat{c}(1) - \bar{c}(1) = \hat{c}(1) - c_{II}(1) = (1 - \sigma)(T_H - T_L) \geq 0. \quad (31)$$

Recall that $\bar{c}(\cdot)$ is linearly increasing in q for $q < q^*$ and linearly decreasing for $q > q^*$ and is thus maximized by q^* . The function $\hat{c}(\cdot)$ is strictly concave. For $\delta = 1$, we have $\hat{c}(q^*) = \bar{c}(q^*)$. This together with (30) and (31) implies that $\hat{c}(q) \geq \bar{c}(q)$ for all $q \in [0, 1]$. Therefore, for $\delta = 1$ the first-best can be implemented by Proposition 3, whenever diagnosis effort is optimal in the first-best.

For $\delta < 1$, by (29) we have $\hat{c}(q^*) < \bar{c}(q^*)$. This implies that $\hat{c}(q) < \bar{c}(q)$ for some values of q sufficiently close to q^* and so diagnosis effort cannot be contractually implemented if $c \in (\hat{c}(q), \bar{c}(q))$, even though it is optimal in the first-best. Q.E.D.

The set of parameter combinations for which the first-best requires a treatment based on information acquisition, but this is not implementable by a contract, is equal to

$$Z \equiv \{c, q | \hat{c}(q) < c < \bar{c}(q)\}. \quad (32)$$

By Proposition 5, this set is non-empty if and only if $\delta < 1$. In Figure 2 the set Z is depicted for this case by the gray shaded area. For a parameter combination in Z only a second-best solution without diagnosis effort can be obtained by the contracting parties. From our previous analysis it follows that the optimal contract then has the following properties:

Proposition 6 *Let $(c, q) \in Z$. Then contractually implementing the high-cost treatment without diagnosis is optimal if $q \leq q^*$, and implementing the trial-&-error strategy without diagnosis is optimal if $q \geq q^*$.*

In comparison with the first–best, the second–best solution involves a higher likelihood of overtreatment for the high-cost treatment, and of undertreatment for the the trial-&-error strategy. More precisely, the following efficiency losses arise

$$\begin{aligned} S_I^* - S_H^* &= (1 - q)\sigma((1 - \delta)(v - T_H) + T_L) - q(1 - \sigma)(T_H - T_L) - c \\ S_I^* - S_{T\&E}^* &= q\sigma(T_H - T_L) - (1 - q)(1 - \sigma)((1 - \delta)(v - T_H) + T_L) - c. \end{aligned} \quad (33)$$

In both expressions, the first term indicates the gain from more precise information, allowing to avoid overtreatment and undertreatment, respectively. The second term arises from the fact that the signal is sometimes incorrect and following it leads to undertreatment and overtreatment, respectively. Lastly, the diagnosis cost has to be taken into account.

5 Separation of Diagnosis and Treatment

Our analysis in the previous sections implicitly assumes that a single expert is responsible for both diagnosis and treatment. We now show that the first–best outcome can be obtained if separating diagnosis and treatment is feasible. As we indicate in the Introduction, this is possible in situations where diagnosis and treatment are essentially independent procedures with small economies of scope. To simplify our analysis, in this section we completely abstract from any kind of interdependencies or economies of scope.³⁰

Suppose the consumer contracts with two different agents, a and b , for diagnosis and treatment. Agent a is an expert for diagnosis and can acquire information about the principal’s problem by investing the effort cost c ; agent b incurs the cost T for providing treatment T . After a diagnosis, agent a prescribes a treatment which is then executed by agent b . Otherwise, the sequence of events and the assumptions on observability are the same as explained in Section 2. A contract specifies the payments

$$p^a \equiv (p_H^a, p_{LS}^a, p_{LF}^a), \quad p^b \equiv (p_H^b, p_{LS}^b, p_{LF}^b) \quad (34)$$

each agent receives, contingent on the first period treatment and the principal’s report about the outcome in case of treatment T_L . As before, when the principal reports failure of treatment T_L , the payment p_{LF}^b includes agent b ’s compensation for the additional treatment T_H in the second period. In total the principal now has to pay

$$p_H \equiv p_H^a + p_H^b, \quad p_{LS} \equiv p_{LS}^a + p_{LS}^b, \quad p_{LF} \equiv p_{LF}^a + p_{LF}^b. \quad (35)$$

To ensure that he reports success and failure of the low–cost treatment truthfully, his total payments have to satisfy condition (15).

³⁰Implicitly we also assume that the agents for diagnosis and treatment are prevented from colluding with each other and exchanging side payments.

After investing the diagnosis cost c and observing a signal $s \in \{s_L, s_H\}$, agent a 's posterior belief that the principal has the minor problem θ_L is equal to

$$\begin{aligned}\pi_L &\equiv \text{Prob}(\theta_L|s_L) = 1 - \text{Prob}(\theta_H|s_L) = \frac{\sigma q}{\sigma q + (1 - \sigma)(1 - q)}, \\ \pi_H &\equiv \text{Prob}(\theta_L|s_H) = 1 - \text{Prob}(\theta_H|s_H) = \frac{(1 - \sigma)q}{(1 - \sigma)q + \sigma(1 - q)},\end{aligned}\tag{36}$$

for signal s_L and s_H , respectively. Note that $\pi_L > \pi_H$ as $\sigma > 1/2$. A contract optimally delegates the choice of treatment to the diagnosis expert a . Thus, after the diagnosis agent a informs agent b about the appropriate treatment.

Since the information obtained by diagnosis effort is not publicly observable, under an optimal contract agent a should truthfully reveal the appropriate treatment that agent b has to execute. The following lemma characterizes the payments p^a that make prescribing the appropriate treatment incentive compatible for agent a :

Lemma 2 *Let (15) hold so that the principal reports success and failure truthfully after treatment T_L . Then agent a prescribes T_H after observing signal s_H , and T_L after signal s_L , if and only if*

$$\pi_H p_{LS}^a + (1 - \pi_H) p_{LF}^a \leq p_H^a \leq \pi_L p_{LS}^a + (1 - \pi_L) p_{LF}^a.\tag{37}$$

Proof: If agent a selects treatment T_H after observing signal s , his payoff is simply p_H^a because this treatment always succeeds. If instead he chooses T_L , his expected payoff after observing signal s_i is $\pi_i p_{LS}^a + (1 - \pi_i) p_{LF}^a$, because the posterior probability of failure is $1 - \pi_i$. Therefore, the first inequality in (37) ensures that choosing T_H after s_H is optimal, and the second inequality that T_L is optimal after s_L . Q.E.D.

Intuitively, the expected payments from prescribing the low-cost treatment must be lower than the payments for the high-cost treatment when agent a believes that the latter is appropriate, and higher when he thinks the low-cost treatment is correct.

Finally, since agent a is employed as diagnosis expert, the contract has to ensure that he invests the information cost c . By doing so he receives the ex ante expected payoff

$$V_I^a(p^a) \equiv q [\sigma p_{LS}^a + (1 - \sigma) p_H^a] + (1 - q) [\sigma p_H^a + (1 - \sigma) p_{LF}^a] - c.\tag{38}$$

Note that the difference with $V_I(\cdot)$ in (19) is that agent a does not incur any treatment costs, because now agent b performs the treatment. When not investing in diagnosis, agent a can either get the payoff $V_H^a(p^a)$ by prescribing T_H or $V_{T\&E}^a(p^a)$ by the trial-&-error strategy, where

$$V_H^a(p^a) \equiv p_H^a, \quad V_{T\&E}^a(p^a) \equiv q p_{LS}^a + (1 - q) p_{LF}^a.\tag{39}$$

Thus, the effort incentive constraint

$$V_I^a(p^a) \geq \max [V_H^a(p^a), V_{T\&E}^a(p^a)] \quad (40)$$

implements diagnosis effort by agent a .

As the choice of treatment is verifiable, agent b can be contractually obliged to provide the treatment prescribed by agent a . Therefore, there are no further incentive problems and a treatment based on diagnosis effort is implemented by the payments (p^a, p^b) whenever the constraints (15), (37), and (40) are satisfied. Actually, in what follows we can ignore constraint (37) because the following lemma shows that it is redundant.

Lemma 3 *Let the effort incentive constraint (40) hold. Then also the treatment incentive constraint (37) is satisfied.*

Proof: Solving the inequality $V_I^a(p^a) \geq V_H^a(p^a)$ for p_H^a yields

$$p_H^a \leq \frac{\sigma q p_{LS}^a + (1 - \sigma)(1 - q)p_{LF}^a - c}{\sigma q + (1 - \sigma)(1 - q)}. \quad (41)$$

Since $c \geq 0$, this implies that the second inequality in (37) holds. Solving the inequality $V_I^a(p^a) \geq V_{T\&E}^a(p^a)$ for p_H^a yields

$$p_H^a \geq \frac{q(1 - \sigma)p_{LS}^a + \sigma(1 - q)p_{LF}^a + c}{(1 - \sigma)q + \sigma(1 - q)}. \quad (42)$$

Since $c \geq 0$, this implies that the first inequality in (37) holds. This proves that (40) implies (37). Q.E.D.

The intuition is simply that investing in costly diagnosis and then ignoring the information cannot be optimal for agent a . In other words, if a contract specified payments that induce information to be ignored, then it would not be optimal to invest in information. By Lemma 3, finding payments so that (15) and (40) hold is sufficient to prove that a treatment based on information acquisition can be implemented. The following proposition shows that this is possible whenever this treatment strategy is first-best.

Proposition 7 *Suppose a treatment based on information acquisition is optimal in the first-best, i.e. $c \leq \bar{c}(q)$. Then the first-best outcome can be contractually implemented by separating diagnosis and treatment with the payments for diagnosis*

$$p_H^a = v - T_H + k^a, \quad p_{LS}^a = v - T_L + k^a, \quad p_{LF}^a = \delta(v - T_H) - T_L + k^a, \quad (43)$$

to agent a , and for treatment

$$p_H^b = T_H + k^b, \quad p_{LS}^b = T_L + k^b, \quad p_{LF}^b = T_L + \delta T_H + (1 - \delta)v + k^b \quad (44)$$

to agent b , where k^a and k^b are some constants.³¹

Proof: From the definition of the first–best surplus of the different treatment strategies in (5), (6), and (7) it immediately follows that agent a 's payoff in (38) and (39) satisfies

$$V_H^a(p^a) = S_H^* + k^a, \quad V_{T\&E}^a(p^a) = S_{T\&E}^* + k^a, \quad V_I^a(p^a) = S_I^* + k^a. \quad (45)$$

under the payments in (43). Therefore, whenever exerting diagnosis effort is optimal in the first–best because $S_I^* \geq \max[S_H^*, S_{T\&E}^*]$, then also agent a 's effort incentive constraint (40) is fulfilled.

The principal's payments in (43) and (44) to both agents sum up to

$$p_H = p_{LS} = p_{LF} = v + k^a + k^b. \quad (46)$$

Therefore, (15) is satisfied.

Q.E.D.

With the payments specified in Proposition 7 the three parties together obtain the joint surplus S_I^* , as in the first–best. At the contracting stage this surplus can be split in an arbitrary way by adjusting the constants k^a and k^b according to market conditions or the parties' bargaining power.

To gain an intuition for the efficiency result, notice that agent b 's incentives can be neglected, because he is contractually obliged to provide the required treatment. Key are the incentives of agent a and the principal. On the one hand, agent a 's payments in (43) give him a share of the first–best surplus and establish thus the correct incentives to invest in effort. On the other, the sum of the principal's payments to the two agents does not depend on the principal's report and establish thus the correct incentives to report honestly.

Is separation of diagnosis and treatment required for this construction? To answer this question consider the framework of Section 4 with a single agent and suppose that the agent is paid following (46), that is $p_H = p_{LS} = p_{LF}$. Clearly, the principal has no incentive to misreport, as (15) is satisfied. Further, the critical cost levels in (21) and (22) can be rewritten as

$$c_A(p, q) = c_I(q) + (1 - \sigma)(1 - q)(1 - \delta)v \quad (47)$$

and

$$c_B(p, q) = c_{II}(q) - \sigma(1 - q)(1 - \delta)v. \quad (48)$$

It is easy to see that if $\delta = 1$, these payments establish the correct effort incentives for implementing diagnosis.³² When $\delta < 1$, however, it is no longer true that $c_A(p, q) = c_I(q)$

³¹These constants have to be chosen such that both the principal and the agents are willing to accept the contract. Their exact values are determined by market conditions or the parties' bargaining power.

³²Note that, as stated above, the payments in Proposition 5 are not unique.

and $c_B(p, q) = c_H(q)$, implying that the effort incentive constraint is distorted. In contrast, separation allows to preserve the correct incentives for agent a and the principal by paying agent b an amount of $(1 - \delta)v$ whenever the principal reports failure. Thus, the role of the payments to agent b is to align the incentives of agent a and the principal, rather than establishing incentives for agent b .

6 Concluding Remarks and Extensions

We have studied a credence good problem in which a consumer relies on the advice of an expert in order to choose one of two services. In our model payments must be designed in order to solve two problems. On one hand, diagnostic effort is costly and there is a moral hazard problem when the expert chooses diagnostic effort. On the other hand, treatment success is not publicly observable and payments must depend on the subjective evaluation of the consumer. We find that such payments can be designed for some range of parameter combinations. This range increases when the discount rate decreases and includes all parameter combinations when the discount rate is zero.

We also show that the first-best is always attainable under the assumption that diagnosis and treatment can be separated at no additional cost. Of course, separation may be inefficient and more costly than joint provision. There could be economies of scope in provision of diagnosis and treatment or there could be costs reflecting the consumer's time lost by consulting several experts.³³ In this situation our analysis indicates that, if the discount rate is positive, there is a trade-off between diagnosis effort incentives and the additional cost of separation. As the discount rate increases, separation becomes more attractive, because the set of parameter combinations for which under combined provision the first-best can be reached shrinks.

Our model generalizes the information technology of the expert that the literature on credence goods usually considers. Further generalizations of that technology are likely to make it more difficult to implement the first-best when the expert provides both diagnosis and treatment. However, our efficiency result when separation of both activities is possible is likely to persist. The basic forces in our model are hence likely to be robust. Consider for instance a setting in which the expert chooses the precision of the signal and the cost of the signal is an increasing and convex function of its quality. The first-best requires that the marginal benefit of higher precision – given by (8) – equals marginal cost and adds an additional constraint that optimal contracts must fulfil. This may make it more difficult to obtain the first-best under joint provision of diagnosis and treatment. When separation is

³³Darby and Karni (1973) footnote 5, Emons (1997, 2001).

possible, however, the agent's payments under the optimal contract differ from the first-best surplus by an additive constant, and set therefore the right incentives.

We have argued that avoiding overtreatment is particularly important in health care, as many countries have experienced large increases in health care spending. In this context it may be reasonable to assume that physicians are to some degree altruistic and value intrinsically providing high-quality care.³⁴ In order to obtain an intuition into the effects these concerns might generate, consider the setting in which the expert provides both diagnosis and treatment. Assume he values in his payoff function the patient's health in some small proportion. Notice that the physician always obtains this additional utility unless the first-period treatment results in undertreatment, in which case this payment is delayed to the next period. As a result treatment options become less attractive the more often they involve undertreatment. In particular, when the likelihood of the minor problem is sufficiently high, the range of parameter combinations for which investment in diagnosis takes place increases as altruism becomes a more important concern. Although this might indicate that with an altruistic expert the moral hazard problem and the problem of subjective evaluation become easier to solve simultaneously, these issues require a more complete analysis. The reason is that whenever the signal is not perfectly precise, there is also (a smaller) countervailing effect. When the likelihood of the minor problem is sufficiently low, prescribing the high-cost treatment becomes preferred to investment in diagnosis, for some parameter combinations. We leave a rigorous study of the optimal contract in the presence of altruism for future research.

³⁴See Green (2014), Liu and Ma (2013) or Henning-Schmidt et al. (2011).

7 References

- ALGER, I. AND F. SALANIÉ (2006): “A Theory of Fraud and Overtreatment in Experts Markets,” *Journal of Economics and Management*, 15(4), 853–881.
- ALONSO, R. AND N. MATOUSCHECK (2008): “Optimal Delegation,” *Review of Economic Studies*, 75(1), 259–293.
- BAKER, G. P., R. GIBBONS AND K. J. MURPHY (1994): “Subjective Performance Measures in Optimal Incentive Contracts,” *Quarterly Journal of Economics* 109, 1125–1156.
- BARON, D. P. (2011): “Credence Attributes, Voluntary Organizations, and Social Pressure,” *Journal of Public Economics*, 95 (11–12), 1331–1338.
- BERWICK, D. M. AND A. D. HACKBARTH (2009): “Eliminating Waste in US Health Care,” *JAMA*, 307(14), 1513–1516.
- BONROY, O., S. LEMARIÉ AND J.-P. TROPÉANO (2013): “Credence Goods, Experts and Risk Aversion,” *Economics Letters*, 120, 464–467.
- BROWN, J. AND D. MINOR (2013): “Misconduct in Credence Good Markets,” mimeo.
- CASWELL, J. A. AND E. M. MOJDUSZKA (1996): “Using Informational Labelling to Influence the Market for Quality in Food Products,” *American Journal of Agricultural Economics*, 78(5), 1248–53.
- CAUSHOLLI, M., W. R. KNECHEL, H. LIN, AND D. E. M. SAPPINGTON (2013): “Competitive Procurement of Auditing Services with Limited Information,” *European Accounting Review*, 22(3), 573–605.
- CLEMENS, J. AND J. D. GOTTLIEB (2014): “Do Physicians’ Financial Incentives Affect Treatment Patterns and Patient Health?,” *American Economic Review*, 104(4) , 1320–1349.
- DARBY, M. R. AND E. KARNI (1973): “Free Competition and the Optimal Amount of Fraud,” *Journal of Law and Economics*, 16(1), 67–88.
- DE JAEGER, K. AND M. JEGERS (2001): “The Physician-Patient Relationship as a Game of Strategic Information Transmission,” *Health Economics*, 10, 651–668.
- DEMSKI, J. AND D. SAPPINGTON (1987): “Delegated Expertise,” *Journal of Accounting Research*, 25(1), 68–89.
- DONALDSON C. AND K. GERARD (1989): “Paying General Practitioners: Shedding Light on the Review of Health Services,” *Journal of the Royal College of General Practitioners* 39(320), 114–117.

- DULLECK, U. AND B.U. WIGGER (2015): "Politicians as Experts, Electoral Control, and Fiscal Restraints," *Journal of Public Economics*, 121, 106–116.
- DULLECK, U. AND R. KERSCHBAMER (2006): "On Doctors, Mechanics, and Computer Specialists: The Economics of Credence Goods," *Journal of Economic Literature*, 44(1), 5–42.
- DULLECK, U. AND R. KERSCHBAMER (2009): "Experts vs. Discounters: Consumer Free-Riding and Experts Withholding Advice in Markets for Credence Goods," *International Journal of Industrial Organization*, 27(1), 15–23.
- DULLECK, U., J. GONG AND J. LI (2013): "Contracting for Infrastructure Projects as Credence Goods," *Journal of Public Economic Theory*, forthcoming.
- EMONS, W. (1997): "Credence Goods and Fraudulent Experts," *Rand Journal of Economics*, 28, 107–119.
- EMONS, W. (2001): "Credence Goods Monopolists," *International Journal of Industrial Organization*, 19, 375–389.
- FEDDERSEN, T.J. AND T.W. GILLIGAN (2001): "Saints and Markets: Activists and the Supply of Credence Goods," *Journal of Economics and Management Strategy*, 10, 149–171.
- FONG, Y.-F. (2005): "When do Experts Cheat and Whom do They Target?," *Rand Journal of Economics*, 36(1), 113–130.
- FONG, Y.-F., L. TING AND D.J. WRIGHT (2014): "On the Role of Verifiability and Commitment in Credence Goods Markets," *International Journal of Industrial Organization*, 37, 118–129.
- FRANKEL, A. AND M. SCHWARZ (2014): "Experts and Their Records," *Economic Inquiry*, 52(1), 56–71.
- FUCHS, V. R. (2009): "Eliminating "Waste" in Health Care," *JAMA*, 302(22), 2481–2482.
- FUCHS, W. (2007): "Contracting with Repeated Moral Hazard and Private Evaluations," *American Economic Review* 97, 1432-1448.
- GABSZEWICZ, J. J. AND J. RESENDE (2012): "Differentiated Credence Goods and Price Competition," *Information Economics and Policy*, 24(3-4), 277–287.
- GIANNAKAS, K. (2002): "Information Asymmetries and Consumption Decisions in Organic Food Product Markets," *Canadian Journal of Agricultural Economics*, 50, 35–50.

- GREEN, E. P. (2014): "Payment Systems in the Healthcare Industry: An Experimental Study of Physician Incentives," *Journal of Economic Behavior and Organization*, 106, 367–378.
- HENNING-SCHMIDT, H., R. SELTEN AND D. WIESEN (2011): "How Payment Systems affect Physicians' Provision Behaviour—An Experimental Investigation," *Journal of Health Economics*, 30, 637–646.
- HOLMSTRÖM, B. (1984): "On the Theory of Delegation," in: *Bayesian Models in Economic Theory*. Ed. by M. Boyer, and R. Kihlstrom. North-Holland, New York.
- HYNDMAN, K. AND S. OZERTURK (2011): "Consumer Information in a Market for Expert Services," *Journal of Economic Behavior and Organization*, 80(3), 628–640.
- JEGERS, M., K. KESTELOOT, D. DE GRAEVE AND W. GILLES (2002): "A Typology for Provider Payment Systems in Health Care," *Health Policy*, 60, 255–273.
- KALE, M. S., T. F. BISHOP, A. D. FEDERMAN AND S. KEYHANI (2011): " "Top 5' Lists Top \$5 Billion," *Archives of Internal Medicine*, 171(20), 1856–1858.
- KNECHEL W.R. (2013): "Do Auditing Standards Matter?," *Current Issues in Auditing*, 7(2), A1–A16.
- LEVIN, J. (2003): "Relational Incentive Contracts," *American Economic Review* 93, 835–857.
- LIU, T. (2011): "Credence Goods Markets with Conscientious and Selfish Experts," *International Economic Review*, 52, 227–244.
- LIU, T. AND C. MA (2013): "Health Insurance, Treatment Plan, and Delegation to Altruistic Physician," *Journal of Economic Behavior & Organization* 85, 79–96.
- MACLEOD, B. (2003): "Optimal Contracting with Subjective Evaluation," *American Economic Review* 93, 216–240.
- MALCOMSON, J.M. (2004): "Health Service Gatekeepers," *Rand Journal of Economics*, 35(2), 401–421.
- MULLAINATHAN, S., M. NOETH AND A. SCHOAR (2012): "The Market for Financial Advice: An Audit Study," NBER Working Paper No. 17929. Available at NBER: <http://www.nber.org/papers/w17929>.
- MYERSON, R. (1979): "Incentive Compatibility and the Bargaining Problem," *Econometrica* 47(1), 61–73.

- MYERSON, R. (1986): "Multistage Games with Communication," *Econometrica* 54(2), 323–58.
- PESENDORFER, W. AND A. WOLINSKY (2003): "Second Opinions and Price Competition: Inefficiency in the Market for Expert Advice," *Review of Economic Studies*, 70(2), 417–437.
- PEARCE, D. G. AND E. STACCHETTI (1998): "The Interaction of Implicit and Explicit Contracts in Repeated Agency," *Games and Economic Behavior* 23, 75–96.
- PITCHIK, C. AND A. SCHOTTER (1987): "Honesty in a Model of Strategic Information," *American Economic Review*, 77, 1032–1036.
- PITCHIK, C. AND A. SCHOTTER (1993): "Information Transmission in Regulated Markets," *Canadian Journal of Economics*, 26, 815–829.
- SÜLZLE, K. AND A. WAMBACH (2005): "Insurance in a Market for Credence Goods," *The Journal of Risk and Insurance*, 72, 159–176.
- TAYLOR, C.R. (1995): "The Economics of Breakdowns, Checkups, and Cures," *Journal of Political Economy*, 103(1), 53–74.
- UNITED STATES GOVERNMENT ACCOUNTABILITY OFFICE (2013): "MEDICARE Higher Use of Costly Prostate Cancer Treatment by Providers Who Self-Refer Warrants Scrutiny," GAO-13-525, available at: <http://www.gao.gov/products/GAO-13-525>, accessed on 23/07/2014.
- WOLINSKY, A. (1993): "Competition in a Market for Informed Experts' Services," *Rand Journal of Economics*, 24, 380–398.
- WOLINSKY, A. (1995): "Competition in Markets for Credence Goods," *Journal of Institutional and Theoretical Economics (JITE) / Zeitschrift für die gesamte Staatswissenschaft*, 151(1), 117–131.