



Review article

Credence goods in the literature: What the past fifteen years have taught us about fraud, incentives, and the role of institutions



Loukas Balafoutas^{a,*}, Rudolf Kerschbamer^b

^a Department of Public Finance, University of Innsbruck, Universitaetsstrasse 15, A-6020 Innsbruck, Austria

^b Department of Economics, University of Innsbruck, Austria

ARTICLE INFO

Article history:

Received 4 February 2020

Accepted 5 February 2020

Available online 10 February 2020

JEL classification:

D82

D83

D21

D22

D18

I11

L15

Keywords:

Credence goods

Expert services

Fraud

Undertreatment

Overtreatment

Overcharging

ABSTRACT

We review the literature on credence goods since Dulleck and Kerschbamer (2006). We consider various markets for credence goods and briefly discuss evidence on the extent of fraud. We then review theoretical and empirical contributions on the determinants of seller and consumer behavior in markets for credence goods. The topics include informational asymmetries, pro-social motivations and seller characteristics, as well as several features of the market structure and institutional environment (separation of diagnosis and treatment, liability, verifiability, reputational concerns, competition between experts and second opinions). We also describe recent developments in this area of research (such as the role of investing in more precise diagnostic technologies) and offer an outlook on future questions.

© 2020 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

Contents

1. Introduction.....	1
2. Prevalence and importance of credence goods markets.....	3
3. Determinants of service provision: Characteristics of consumers and expert sellers.....	4
3.1. Consumer information.....	4
3.2. Non-monetary incentives and motivations of expert sellers.....	5
4. Determinants of service provision: Institutional and market characteristics.....	6
4.1. Separation of diagnosis and treatment.....	6
4.2. The role of insurance.....	8
4.3. Liability and verifiability.....	9
4.4. Competition between sellers.....	10
4.5. Reputation.....	11
4.6. Second opinions.....	11
5. Additional dimensions of the seller–consumer interaction: Diagnostic effort, prices, and tax attitudes.....	13
6. Concluding remarks.....	14
References.....	15

1. Introduction

What do doctors, lawyers, taxi drivers, car mechanics, computer specialists and financial advisors have in common? From an economic perspective, a common characteristic that stands

* Corresponding author.

E-mail address: loukas.balafoutas@uibk.ac.at (L. Balafoutas).

out is the superior information those experts possess vis-à-vis the consumers they interact with. This is true *ex ante*, since the expert provider can typically observe the kind or quality of the good or service the consumer needs while the consumer cannot; but it is often also true *ex post*, since the consumer can often not even after consumption verify the type or quality of the service that she has received and paid for.¹ Markets such as these, in which informed expert sellers interact with uninformed consumers, are known as markets for *credence goods* (Darby and Karni, 1973), or markets for expert services. As almost any market characterized by asymmetric information, such markets are prone to inefficiencies.

The early literature on credence goods did not yield a clear picture regarding the inefficiencies arising from the informational asymmetries and about possible remedies. Rather, the results and policy implications seemed to depend sensitively on the specific conditions of the settings considered. Moreover, apparently similar models often led to contradicting conclusions. Almost fifteen years ago, Dulleck and Kerschbamer (2006) – henceforth DK06 – presented a simple unifying model of credence goods and a small set of organizing assumptions with the property that the price mechanism alone is sufficient to solve the fraudulent expert problem if those assumptions are satisfied. They showed that most existing results on inefficiencies and fraud in credence goods markets can be reproduced by dropping one of those assumptions.

Specifically, DK06 identify two basic problems on markets for credence goods. First, the provision of goods or services that do not fit the needs of the customers, and second, the charging for goods or services that have not been provided. For the former problem the authors distinguish between two subcategories: Underprovision or *undertreatment* refers to the case where the service provided by the expert seller is insufficient to meet the needs of the customer. Consider, for instance, the case of a broken cooler in a car: If the mechanic replaces only the gasket then the cooler might work for some days but might then again start to trickle. If it does, then we would qualify the service provided by the mechanic as underprovision because the repair was insufficient to solve the problem. Underprovision is per definition inefficient, since there is a resource cost but (almost) no benefit. The opposite case, where the consumer requires a simple and inexpensive service but receives a sophisticated one is called overprovision, *overtreatment*, or supplier-induced demand. Remaining with the car repair example, we speak of overprovision if the mechanic replaces the cooler although replacing the gasket would have solved the problem. Overtreatment is also inefficient because the additional benefits to the consumer from the sophisticated repair are smaller than the additional costs. The second problem mentioned earlier – putting a more expensive service on the bill than the one actually provided – is termed *overcharging*. In our car repair example the mechanic might charge for a new cooler although only a new gasket was provided. In the short run, overcharging is a pure transfer from the consumer to the expert. In the long run, it might also lead to inefficiencies if the fear of becoming the victim of overcharging induces customers to search for multiple opinions, to postpone repairs or even to leave the market.

In their simple unifying model of credence goods DK06 identify the following conditions under which the market eliminates the incentives for expert dishonesty: (i) consumer *homogeneity* in the likelihood of having one of two different types of problem and in their valuation for a solution of that problem; (ii) *commitment* by the consumer to accept the service recommended by the expert seller once she has received a recommendation,

and commitment by the expert seller to provide a service at the price he has posted *ex ante* for the service; and (iii) either *liability*, which means that the expert seller cannot provide a cheap service when a sophisticated one is needed, or *verifiability*, which means that the seller cannot provide a cheap and charge for an extensive service.

Liability rules solve the undertreatment problem (trivially) but they do not directly address the other two problems (overtreatment and overcharging). DK06 show that under *liability* (but without verifiability) the prices of the different services tend to converge to a ‘uniform price’. Under a uniform price an expert has neither an incentive to change the cooler when only a new gasket is needed (this causes additional costs for the mechanic without leading to a higher price), nor an incentive to charge for a new cooler when only the gasket has been replaced (because the prices for the two services are the same). Without liability, the market should still work efficiently if the consumers can verify which service has been provided. Under *verifiability* (but without liability) the overcharging problem is eliminated more or less by definition, but the problems of under- and overtreatment are still present. DK06 show that the market will solve those problems by leading to equal mark-up prices: If the profit of the mechanic does not depend on whether he sells the new cooler or only a new gasket, then he has neither an incentive to sell a new gasket to a car owner who needs a new cooler nor an incentive to sell a new cooler to a car owner who needs only a new gasket. Taking the efficiency results derived under the above mentioned conditions as the starting point, DK06 then describe in detail the nature of fraud and inefficiencies that arise when (at least some) of these conditions are absent. Thereby they summarize much of the earlier literature on this topic.

In this paper we review the literature on credence goods since DK06. The number of theoretical and empirical contributions has continuously increased during this time, leading to a large and flourishing body of literature on this empirically highly relevant type of market. In order to keep the scope of this review tractable, we focus on papers with an explicit relation to the problem and nature of credence goods. This means that we place little focus on discussing issues in relevant markets (such as healthcare) other than the ones arising from the informational asymmetries outlined above. Similarly, we refrain from a more general discussion of inefficiencies in the markets under consideration and concentrate on the dimensions of mistreatment and fraudulent charges by expert sellers.²

Our review is structured as follows. In Section 2 we begin by briefly discussing various markets for credence goods, as well as giving the reader an idea of the extent of fraud as documented in several empirical studies. In Sections 3 and 4 we describe a number of factors that are likely determinants of behavior and outcomes in markets for credence goods. Section 3 focuses on two of the most important characteristics of buyers and sellers – the information of buyers and the pro-social motivation of sellers, respectively. Section 4 considers aspects of the market structure and institutional environment – the possibility of separation between diagnosis and treatment, the role of liability and verifiability, reputational concerns and competition between expert sellers, and the potential impact of searching for multiple opinions. We believe that a careful consideration of all of these determinants of market performance is of particular value, since it can be used as a guide for policy recommendations to devise fraud-restraining institutions and to improve efficiency. In Section 5 we turn to some other important dimensions of the

² An exception is Section 5 where we also discuss the moral hazard problem that arises if the expert must exert costly but unobservable effort to identify the service that meets the consumer’s need.

¹ In this article, the expert seller is male and his customer is female.

seller–consumer interaction that are not included in the previous sections, namely diagnostic effort, pricing, and tax attitudes.³ Finally, Section 6 concludes and also offers some thoughts on the future of research on credence goods.

2. Prevalence and importance of credence goods markets

In this section we briefly review some studies that present evidence from various sectors, either from audit studies or using existing field data, to underline the pervasiveness of fraud in markets for expert services. The economic importance and relevance of credence goods has already been documented by a number of studies from the 1990s which provide evidence from the field suggesting that expert sellers exploit their informational advantage in order to increase their profits at the expense of uninformed consumers.

The earlier literature in this domain comes predominantly from the *healthcare sector* and is – in part – already discussed in DK06. For the sake of exposition, we mention here [Domenighetti et al. \(1993\)](#) who show that better information – in the form of having at least one doctor in one's family – is associated with a substantially lower likelihood of receiving surgical treatment, and [Gruber and Owings \(1994\)](#) who show that the proportion of caesarian deliveries responds to the fee differential between caesarian and natural childbirth. Similarly, a study with Japanese data has found similar effects with respect to drug prescriptions and how these respond to mark-ups ([Iizuka, 2007](#)). More recent studies lend further support to the pervasive nature of fraud by expert sellers in healthcare services. [Baker \(2010\)](#) analyzes U.S. data from Medicare for the period 1999–2005 and reports a large increase in the number of ordered MRI scans, following the acquisition of MRI equipment by non-radiologist physicians who could bill for the scans themselves (instead of referring patients to hospitals or other specialized facilities). [Clemens and Gottlieb \(2014\)](#) use data from Medicare in the US and find a positive relationship between physician payments and the supply of medical services, in particular those with a large discretionary power on part of the physician. This suggests overtreatment by physicians in the presence of strong financial incentives. A similar effect, namely a positive correlation between prices and the extent of healthcare services provided, is documented for India by [Das et al. \(2016\)](#).

[Gottschalk et al. \(2020\)](#) examine in a field experiment how overtreatment in the Swiss market for *dental care* depends on a number of patient, dentist, and market characteristics. For the purpose of their experiment one patient visits 180 different dentists and records the received recommendations. Given that the patient needs no dental treatment, any recommendation for treatment is classified as unnecessary and hence as fraudulent. The overall overtreatment rate is 27.8% (50 of 180 cases), and it is significantly higher for patients of low – as opposed to high socio-economic status.

A number of studies have arrived at similar results in sectors other than healthcare. For the *car repair* sector, [Wolinsky \(1995\)](#) cites the estimate of the US Department of Transportation indicating that about one half of all car repairs undertaken in the US is unnecessary. In line with this estimate, [Schneider \(2012\)](#) finds in a natural field experiment that more than half of the car repairs undertaken during undercover visits to car mechanics were in fact unnecessary. Another study relating to the car repair

sector is the lab experiment by [Beck et al. \(2014\)](#) who use a non-standard population of professional car mechanics and find that car mechanics are significantly more prone to supplying unnecessary services than other subjects, which could be a result of decision heuristics they learned in their professional training.

In the market for *taxi services*, [Balafoutas et al. \(2013, 2017\)](#) quantify the extent of fraudulent behavior by taxi drivers in the form of overtreatment (taking unnecessary detours in order to increase the total fare) and overcharging (artificially increasing the fare by charging fake additional charges or manipulating the taximeter). [Balafoutas et al. \(2013\)](#) report that about 45% of the taxi drivers in their sample took at least a short detour, while 11% engaged in overcharging, requesting a higher price than the one corresponding to the distance actually driven.

Several natural field experiments document the prevalence of fraud in the market for *computer repairs*. [Kerschbamer et al. \(2016\)](#) investigate how informing a repair shop that an insurance agency will cover the repair costs affects the provision and charging behavior of computer repair shops in Austria. The authors find robust evidence for fraud, mostly in the form of overcharging in the working time dimension, but also some evidence for overtreatment (replacing parts that did not need replacement). Similar evidence is reported by [Kerschbamer et al. \(2019\)](#) for the German market for computer repairs.

A promising and exciting recent development is the use of the credence goods framework for the study of fraud in the market for *financial advice*, mainly in the form of mistreatment.⁴ The idea here is that advisors are experts who typically possess more information than their clients about the fit between the characteristic of an asset and the preferences of the client, but who might have misaligned incentives because their income consists in parts of commissions. A series of interesting theoretical studies in this line of literature are the ones by [Inderst and Ottaviani \(2012a,b,c\)](#), which acknowledge that advisors often act as brokers and are driven by financial incentives based on commissions paid to them by product providers. These studies are preceded by [Inderst and Ottaviani \(2009\)](#), a paper that sets the stage for the analysis of financial advice as a credence good by modeling mis-selling by salespersons who act as intermediaries (brokers) of the product provider. [Inderst and Ottaviani \(2009\)](#) focus on various aspects of the agency relationship between the selling firm and its brokers and on how these aspects affect the suitability of recommendations to the final consumers. [Mullainathan et al. \(2012\)](#) present empirical evidence that goes in the same direction as the predictions of the studies by Inderst and Ottaviani: in an audit study with undercover clients they find that financial advisors generally prefer to recommend financial products that are aligned with their own personal interests (e.g., products with high commissions).

[Anagol et al. \(2017\)](#) present evidence from an audit study on *insurance advice* in India, in which several auditors seek advice on life insurance products from insurance agents. One of the key findings of this study is that insurance agents commonly recommend products that do not cater to the consumers' needs, but rather increase commissions. The rate at which insurance agents recommend suitable products is found to be as low as 5% in this setting.

For completeness, it is also worth referring to two recent studies that consider *politicians in the role of experts*. In [Dulleck et al. \(2013\)](#) politicians are experts who are better informed than the electorate about public infrastructure projects, while

³ It should be noted that we have opted to structure our review along topics rather than along methodology. This means that, for each of the above topics, our discussion includes both theoretical and empirical (mostly experimental) studies. This clearly distinguishes the current paper from [Kerschbamer and Sutter \(2017\)](#) that focuses almost exclusively on field experiments in markets for credence goods.

⁴ The terms *undertreatment* and *overtreatment* are well defined only in a vertically differentiated market. In a horizontally differentiated market we speak of *mistreatment* when the service provided does not fit to the needs of the consumer.

in [Dulleck and Wigger \(2015\)](#) the superior information refers to budgetary issues and the amount of public spending that is needed to maximize economic welfare. While these studies approach the problem from a theoretical perspective and do not provide empirical evidence, they are novel as they are the first to cast the relationship between politicians and the electorate into a credence goods context.

3. Determinants of service provision: Characteristics of consumers and expert sellers

3.1. Consumer information

Given that the defining feature of a credence good is an informational asymmetry between the seller and the buyer, it is natural to presume that the outcome on such a market is sensitive to changes in the degree of this asymmetry. It is therefore not surprising that several theoretical studies have focused on the relationship between consumer information and fraud by expert sellers.

[Hyndman and Ozerturk \(2011\)](#) investigate a credence goods model in which the consumer's problem is either minor (requiring a cheap treatment) or major (requiring an expensive treatment). Regarding verifiability and liability the authors assume that the treatment provided is not observable by the consumer (implying that overtreatment is dominated by overcharging) and that the expert is liable to solve the customer's problem (implying that undertreatment is ruled out). Thus, the focus of the paper is on the overcharging dimension of fraud. Departing from DK06, the authors assume that consumers are heterogeneous in two dimensions, in their valuation for getting the major problem fixed and in their information regarding the problem they have. Regarding information they assume that, prior to visiting the expert, some consumers receive a noisy signal about whether they have a minor or a major problem. In this framework there is a cheating equilibrium, in which low-valuation consumers are never cheated, while high-valuation consumers are cheated with positive probability. Interestingly, in this equilibrium more information can be good or bad for a high-valuation consumer, depending on what the information reveals about the type of problem she has: Among high-valuation consumers, the most likely victims of cheating are those who receive a signal indicating that their problem is likely to be major, followed by the uninformed consumers and finally those who received a good signal.

The intuition for this result is as follows: The cheating equilibrium the authors characterize is a mixed strategy equilibrium very similar to the one discussed in Lemma 6 of DK06. In equilibrium the expert does not overcharge a high-valuation consumer with the minor problem all the time, but only with a strictly positive probability because a cheap recommendation is always accepted while an expensive recommendation is only accepted with some probability. To induce some high-valuation consumers to accept and others to reject the expensive recommendation, consumers must be indifferent between accepting and rejecting the expensive recommendation. Since consumers who have received a bad signal are more likely to have the major problem and therefore more willing to accept the expensive recommendation, they have to be cheated more frequently to keep them indifferent between acceptance and rejection.

[De Jaegher \(2012\)](#) investigates a very similar framework, the main differences to [Hyndman and Ozerturk \(2011\)](#) being that (i) the expert does not observe the consumer's information, and (ii) the treatment provided is verifiable. The former difference implies that the expert cannot condition his behavior on the consumer's information in [De Jaegher \(2012\)](#), while such a dependence is one of the drivers of the results of [Hyndman and](#)

[Ozerturk \(2011\)](#). The latter difference means that cheating comes in the form of overtreatment in [De Jaegher \(2012\)](#), while it involves overcharging in [Hyndman and Ozerturk \(2011\)](#). For the case where the expert has strong incentives for overtreatment, [De Jaegher \(2012\)](#) characterizes a mixed strategy equilibrium where a consumer with a bad signal is indifferent between accepting and rejecting the expensive recommendation, whereas a consumer with the good signal rejects with certainty. In this equilibrium more information hurts the consumer with the bad signal, because she is overtreated more often in order to keep her indifferent (as a more informative bad signal would otherwise tilt the balance towards acceptance).

The idea that more information is not necessarily beneficial for consumers in a market for credence goods is also found in the model of [Inderst and Ottaviani \(2012a\)](#). The authors evaluate the effects of mandatory disclosure of commissions (aimed at improving consumers' information and educating them about potential conflicts of interest on the side of advisors) and show that such a policy can have unintended efficiency consequences by reducing the market share of more cost-efficient firms. The paper by [Loewenstein et al. \(2011\)](#) also suggests that disclosure can backfire, by leading to more biased advice and at the same time reducing the likelihood that such advice is turned down by consumers.

Turning to empirical work, [Schneider et al. \(2016\)](#) offer evidence from the economic lab on the effects of enhanced consumer information on cheating behavior by experts. The theoretical framework for the experiment builds on [Hyndman and Ozerturk \(2011\)](#) who (as explained above) predict that the most likely victims of cheating are high-valuation consumers who receive the signal that the problem is likely to be major. The results in [Schneider et al. \(2016\)](#) are in line with this prediction: Experts in the lab indeed condition their cheating behavior on the consumer's risk of suffering from a serious problem if they can observe her information. Accordingly, experts and low-risk consumers benefit from the additional information given to consumers – at the expense of the more frequently cheated high-risk consumers.

A number of field experiments relate fraud by expert sellers to consumer information. An example is [Balafoutas et al. \(2013\)](#) who systematically investigate the impact of varying degrees of consumer information on expert sellers' behavior. In the market for taxi rides in Athens, Greece, the authors collected 348 observations using research assistants as undercover passengers. These passengers belonged to one of three 'information roles' – local, nonlocal native, and foreigner, respectively. Passengers in all three information roles instructed the driver upon entering the taxi to take them to a particular destination. Passengers in the roles of locals and non-local natives did this in Greek, whereas passengers in the role of foreigners spoke in English. Passengers in the role of non-local natives and of foreigners then asked the driver whether he knew the destination, adding as an explanation for asking that they were not familiar with the city. Given that non-local native and foreign passengers signaled to the driver that they were unfamiliar with the city, whereas local passengers did not, and given that theory predicts that there is not much room for overtreatment of passengers who know the shortest route to their destination, the former two groups are predicted to be more prone to be taken on detours than the latter group. Since an English-speaking passenger is arguably less likely perceived to be familiar with the details of the Greek taxi tariff system, and since theory predicts that there is not much room for overcharging of consumers who can verify whether the correct tariff has been applied, passengers in the role of foreigners are predicted to be more prone to overcharging than passengers in the other two information roles. The main finding in [Balafoutas et al. \(2013\)](#) is that fraudulent behavior by taxi drivers is indeed responsive to

passenger information as perceived by the driver: Passengers in the role of non-local natives and of foreigners were more prone to overtreatment and were taken on detours of about double length compared to local passengers. In the charging dimension, the authors report that creating the impression of being unfamiliar with the tariff system strongly increases the likelihood of fraud in the market under consideration, from 6% to 22%.⁵

Evidence pointing in the same direction is provided by Anagol et al. (2017) for the life insurance industry. In a series of field experiments investigating the quality of advice provided by life insurance agents in India, the authors manipulate in one of their treatments the agent's perception about the customer's level of sophistication about life insurance policies. The authors find that consumers who appear to be more sophisticated receive better advice from insurance sellers.

Kerschbamer et al. (2019) examine in a natural field experiment in the market for computer repairs whether consumers benefit from using online sources to inform themselves before purchasing credence goods. In one of the waves of their experiment, secret shoppers utter either no conjecture about the problem, a correct conjecture, or a false conjecture. Price distributions are indistinguishable in the first two treatments, while in the third prices double. Since the diagnosis provided by specialized software on the internet is almost always noisy and since consumers cannot distinguish between a correct and an incorrect diagnosis, the authors conclude that if a consumer is not absolutely certain about the nature of a problem, she is better off keeping silent.

Finally, in the finance industry, an interesting recent study by Bucher-Koenen and Koenen (2015) investigates the relationship between a client's financial literacy and the quality of financial advice she receives. The authors provide a simple analytical framework in which better informed consumers have a more valuable outside option because they are able to obtain a better alternative on their own. Taking this into account, the advisor has an incentive to provide better advice to consumers who appear to be smarter. This theoretical prediction is then tested empirically using a large dataset which is representative for German households. The evidence presented by the authors indicates that individuals who appear to have lower financial literacy (in particular, subjects without tertiary education and female subjects) do indeed receive worse financial advice.

3.2. Non-monetary incentives and motivations of expert sellers

A number of papers consider the possibility that at least a fraction of experts on credence goods markets displays pro-social preferences.⁶ Liu (2011) models this situation theoretically and arrives at the somewhat counter-intuitive conclusion that the presence of conscientious experts may actually result in more

fraud compared to a situation in which all experts are selfish. In the model of Liu the consumer's problem is either minor or major, the consumer's valuation for a solved major problem is larger than the valuation for a solved minor problem, and there is liability but no verifiability. Under this constellation the monopolistic expert posts a uniform price equal to the expected valuation of the consumer, which the latter accepts. The market outcome then is fully efficient and entails no fraud. Like in earlier work by Fong (2005), in Liu's model this solution is infeasible with only selfish experts in the market because the expected valuation of the consumer is assumed to be lower than the cost of the major treatment and because a selfish expert cannot commit ex ante to treat the consumer at a price below cost. For the case with only selfish experts, Liu (2011) characterizes a no cheating equilibrium where the expert posts two prices, one for the cheap and one for the expensive treatment, and he recommends the cheap treatment for the minor problem and the expensive treatment for the major problem. In equilibrium the expert does not overcharge a consumer with the minor problem, because a cheap recommendation is always accepted while an expensive recommendation is only accepted with some probability that is small enough to make overcharging unprofitable.

The main innovation of Liu (2011) is the introduction of conscientious experts in the market. A conscientious expert derives utility from the material payoff and from repairing the consumer's problem. If the weight on the repair-utility in the conscientious expert's profit function is large enough, then he is able to commit to solving the consumer's problem at a uniform price below the cost of the expensive repair. With only conscientious experts in the market, the profit maximizing uniform price is equal to the expected valuation of the consumer and the equilibrium is fully efficient just like in the full commitment case investigated by DK06. With conscientious and selfish experts in the market, there exist pooling equilibria in which both types of expert post the same uniform price. The uniform price results in a positive profit for the conscientious expert when the problem is minor and in a loss when the problem is major, and since his aversion against leaving problems unrepaired is strong, he repairs both problems. By contrast, the selfish expert repairs only the minor problem but rejects the major one to avoid a loss. Since the uniform price in the pooling equilibrium is above the valuation of the consumer for the cheap treatment, the author concludes that '*when the selfish expert treats the minor problem, he overcharges the consumer*' (p.228). Based on this somewhat unusual definition of overcharging, the author concludes that the presence of a conscientious expert may result in more fraudulent behavior by the selfish expert. Regarding efficiency, there exist parameter constellations in which the social loss in the (most profitable) uniform-price equilibrium in the presence of a conscientious expert is larger than the social loss in the (most profitable) mixed strategy equilibrium in the absence of conscientious experts. In this sense the presence of conscientious experts might make things worse and not better.

In a similar vein, Fong et al. (2014) outline a model in which the presence of some honest experts decreases the range of parameter constellations in which the selfish expert serves the customer efficiently. Important differences between this model and Liu (2011) are, (i) that there is liability but no verifiability in Liu (2011) and verifiability but no liability in Fong et al. (2014), and (ii) that the conscientious expert in Liu (2011) maximizes a weighted sum of her profit and the customer's valuation for a successful intervention, while the honest expert in Fong et al. (2014) makes honest recommendations regardless of his profit from doing so. The latter difference implies that the conscientious expert in Liu (2011) recommends strategically to maximize his

⁵ A seemingly related study is Currie et al. (2011) who conduct an audit study in China in which they send pairs of patients to the same physician and record whether the physician has recommended use of antibiotics. The study varies the perceived degree of patient information by having one of the patients reveal some knowledge about antibiotics, and finds that patients who signal a higher degree of information are less likely to be prescribed antibiotics. This treatment variation bears similarities with the treatment variation in Balafoutas et al. (2013). Note, however, that in this study the variation in patient information does not involve the primary informational asymmetry in a credence goods market (i.e., information about the fit between the property of the good and the characteristic of the consumer) but rather the property of the good per se.

⁶ In the model of Inderst and Ottaviani (2012a) the advisor cares directly that the consumer purchases the product that best fits her needs, among other reasons out of ethical and professional concerns for the consumer's well-being. However, the focus of that paper is mainly on the role of competition between product providers through commissions for advisors and on the effect of disclosure about the level of commissions.

payoff, while the honest expert in Fong et al. (2014) is a behavioral type who is constrained to make honest recommendations regardless of the material incentives.

The starting point of Fong et al. (2014) is the observation that the standard solution to the credence goods problem under verifiability but no liability – equal mark-up prices that provide no material incentives for dishonest recommendations – might not be feasible if the client has the freedom to reject recommendations *ex post*.⁷ In this case the expert is constrained to post price vectors that have the property that the customer has a non-negative net surplus for each treatment. As a consequence, the expert faces a trade-off between surplus maximization and rent extraction and this tradeoff results in over- or undertreatment for some parameter ranges. The authors show that if the surplus from solving the major problem is larger (smaller) than that of solving the minor problem and the major (minor) problem is sufficiently likely, the highest profit is achieved in an equilibrium involving full overtreatment (full undertreatment). On the other hand, when the problem associated with the larger surplus occurs with a smaller probability, the highest profit is achieved by an equilibrium involving partial mistreatment or an equilibrium involving honest and efficient treatment.

In this framework, Fong et al. (2014) examine the role of honest experts and find that the existence of honest experts increases the parameter range in which the opportunistic expert recommends dishonestly and treats inefficiently. The intuition is that in an over- or undertreatment equilibrium the presence of honest experts increases the customer's willingness to accept treatment recommendations which would have been rejected otherwise. This allows the opportunistic expert to raise his price and make more profit. Since this profit-increasing effect is only present in an inefficient but not in the efficient equilibrium, the range of parameter constellations in which an inefficient equilibrium prevails becomes larger.

Beck et al. (2013) study the role of guilt aversion for the provision and charging behavior of expert sellers in theory and then test some core predictions in a laboratory experiment. In line with the theory of guilt aversion, which predicts that a decision maker will suffer a cost if he fails to live up to the perceived expectations of an interaction partner, the authors find that experts who have made a non-binding promise to a consumer are indeed less likely to engage in fraud. Such promises may refer either to the provision decision, the pricing decision, or both, and they are found to increase interaction rates by consumers and to reduce undertreatment and overcharging by experts.

A recent study explicitly geared towards understanding the role of social preferences on credence goods markets is Kerschbamer et al. (2017). This paper starts from the hypothesis that humans are often motivated by forces beyond material self-interest and shows theoretically and verifies in lab experiments that non-trivial 'distributional preferences' of sellers (like altruism, inequality aversion, spite, or envy) can explain the reaction of experimental sellers to changes in the institutional framework much better than the standard assumption of selfish preferences (see Section 4.3 for further details). One important difference between this study and the ones discussed in this section so

far is that Kerschbamer et al. (2017) consider the possibility of heterogeneous non-selfish preferences by sellers, while the other papers assume that all pro-social sellers are of one (conscientious, honest, or guilt averse) type. An important insight from this study is that, instead of choosing doctors, mechanics or computer specialists exclusively according to their training, customers on markets for credence goods should worry more about the attitudes of these experts towards their customers.

In their experimental study on payment schemes for physicians, Hennig-Schmidt et al. (2011) provide evidence indicating that financial incentives are not the only motivating factor for physicians, who seem to be driven also by a desire to serve their patients' needs and adjust their treatment decisions in order to increase health benefits. This finding is confirmed in the lab experiments of Godager and Wiesen (2013) who report that physicians are not solely driven by financial incentives but exhibit some degree of altruism – albeit with a substantial degree of heterogeneity between physicians – and of Hennig-Schmidt and Wiesen (2014) who compare medical and non-medical students and find significantly more other-regarding behavior towards patients in the former sample. Still in the context of payment schemes for physicians, the experiment by Green (2014) reports that intrinsic motivation is in fact the primary motive of physicians and raises a word of caution: retrospective payment schemes have the potential to deteriorate health outcomes because they can crowd out the intrinsic motivation of conscientious physicians.⁸

4. Determinants of service provision: Institutional and market characteristics

4.1. Separation of diagnosis and treatment

In many of the theoretical models of credence goods markets discussed so far (DK06; Hyndman and Ozerturk, 2011; Liu, 2011; De Jaegher, 2012; Fong et al., 2014), the expert seller who possesses or acquires private information about a consumer's problem is also the one who recommends and subsequently provides a treatment to the consumer.⁹ In this subsection we discuss the possibility that, for instance due to legal or institutional constraints, the diagnosis and recommendation on the one hand and the treatment provision on the other hand are provided by two different agents. Intuitively, this could eliminate the incentives of experts to provide an inappropriate treatment recommendation, since they will not reap any financial benefits from selling the recommended treatment (unless they are able to collude with the agent who provides it).

Separation was discussed as a possible solution to the credence goods problem already by Darby and Karni (1973), who introduced this kind of good to the literature. The idea is that the incentive to misreport the diagnosis outcome disappears if the expert cannot profit from the sale of the good. DK06 relate the efficiency consequences of this solution to the magnitude of economies of scope between diagnosis and treatment: For markets where such economies are low (an example is the prescription and sale of drugs) separation is likely to be an efficient solution to fraud by expert sellers, while in markets where those

⁷ To see the problem, assume for the moment that there is a monopolistic expert and a single consumer who is committed to accept the recommendation *ex post* if she was willing to undergo a diagnosis *ex ante*. Under this condition the expert will post a price pair specifying equal mark-ups for both treatments and these mark-ups will be such that the client has a negative surplus from one treatment and a positive surplus from the other, rendering her just indifferent between visiting the expert and not visiting the expert. This solution is no longer feasible if the client cannot commit *ex ante* to accept the treatment *ex post* because it induces the client to reject the recommendation if it implies a negative surplus.

⁸ A large literature in health economics going back to Arrow (1963) adopts the assumption that physicians are motivated by altruistic concerns (see, for instance, Chalkley and Malcomson, 1998; Jack, 2005; Liu and Ma, 2013; Makris and Siciliani, 2013). Given that this strand of the literature does not consider the physician-patient interaction within the framework of a credence goods problem, we do not describe it in detail here.

⁹ A notable exception are the studies by Inderst and Ottaviani (2009, 2012a,b,c), where product providers sell their products to final consumers through brokers, or advisors.

economies are large (as for instance, in the car repair market or in surgery) separation is more likely to simply replace one inefficiency by another.

While the point that the incentive problem of revealing the diagnosis outcome can be solved by breaking up the joint provision of diagnosis and treatment seems rather obvious, a recent paper by [Bester and Dahm \(2018\)](#) makes the more intriguing point that separation might help to obtain the first-best outcome even in a more complicated world where diagnosis requires effort and where neither the diagnosis effort nor the diagnosis outcome are observable. In their model the consumer has again one of two problems – minor or major – and the expert can again provide one of two treatments – cheap or expensive. Departing from much of the earlier literature, the authors assume that the expert has to invest in diagnosis to obtain a noisy signal about whether the cheap treatment is sufficient or not. Performing the diagnosis is costly for the expert and the diagnosis effort is not observable for the consumer. There is verifiability, but only the consumer observes success or failure of the cheap treatment. Payments can therefore depend on the treatment provided and, for the case where the cheap treatment has been provided, also on the consumer's report on treatment success. The authors focus on contracts that specify one payment (or price) for the expensive treatment, one for the cheap treatment when the consumer reports success, and one for the cheap treatment when the consumer reports failure. If the consumer reports failure of the cheap treatment then the expert is obliged to provide the expensive treatment for free in the next period. An optimal contract has to address three problems in this model: (i) a subjective evaluation problem on the consumer side: the consumer should have no incentive to misreport success of the cheap treatment as failure or vice versa; (ii) a moral hazard problem on the expert side: for parameter constellations where efficient service requires diagnosis, the contract should induce the expert to perform the costly diagnosis; (iii) a truthful reporting problem on the expert side: the contract should not provide incentives for the expert to recommend the cheap treatment when the private signal indicates that the expensive treatment is required, and vice versa.

The authors show that the first-best solution can always be obtained in this framework if the discount factor is one and that lowering the discount factor reduces the range of parameter combinations for which the first-best can be obtained.¹⁰ Specifically, the first-best is still implementable for parameter constellations where it dictates provision of the low- or the high-cost service without diagnosis, but it is infeasible for constellations where the first-best entails a costly diagnosis. The problem here is that with a discount factor strictly smaller than one, part of the resource costs of failure are borne by the consumer: Instead of getting her problem solved immediately, she must wait one period, resulting in a utility loss. In this case, any misbehavior by the expert potentially imposes a negative externality on the consumer. Taking this inefficiency result as the starting point, [Bester and Dahm \(2018\)](#) show that the first-best can always be obtained if diagnosis and treatment can be separated by contracting with two different

experts – a diagnosis expert and a provision expert. The intuition is that only the diagnosis expert needs to be incentivized, so the provision expert can be taken as the third party that allows breaking the budget constraint. Specifically, the payments from the consumer to the two experts can be structured in such a way that (i) the sum of the payments to the two experts for the cheap treatment is independent of the report of the consumer, and (ii) the payment the diagnosis expert receives for the cheap treatment is lower if the consumer reports failure than for the case where she reports success. This is achieved by making the payment to the provision expert for the cheap treatment higher for the case where the consumer reports failure and lower for the case where she reports success.

Besides the literature discussing explicitly the role of separation between recommendation and treatment as a solution to the credence goods problem, there is a more applied literature in the healthcare context that takes separation as given and studies the role of different payment schemes for the expert for his behavior. Broadly speaking, one defining feature of payment schemes for physicians is whether the payment includes a variable component (as opposed to a fixed wage) and, if so, how this variable component is determined. For instance, one has to differentiate between fee-for-service systems (the payment that a physician receives per patient depends on the treatment that the patient receives) and capitation systems (the physician receives a fixed, pre-determined fee per patient treated).

[Hennig-Schmidt et al. \(2011\)](#) conduct a lab experiment studying the influence of fee-for-service versus capitation systems on the supply of (verifiable) healthcare services and health benefits received by patients. The participants in the experiment are medical students who are asked to choose a treatment quantity for a patient – that is, the experiment is framed as a medical decision. Consistent with physicians' financial incentives, the supply of medical services is higher under fee-for-service than under capitation. In general, patients tend to be subject to overtreatment under fee-for-service and to undertreatment under capitation. [Brosig-Koch et al. \(2016\)](#) confirm the overtreatment result of [Hennig-Schmidt et al. \(2011\)](#) by showing that physicians, medical students, and non-medical students provide significantly more medical services under fee-for-service than under capitation.

[Green \(2014\)](#) compares the provision behavior of physicians and the health benefits of patients in a framed real-effort lab experiment under several prominent physician payment structures including fee-for-service, capitation, salary, and pay-for-performance. Under the latter the payment to the expert seller is in part contingent on some pre-determined quality measure such that improved health outcomes for patients lead to higher payments for physicians. Pay-for-performance is combined with a capitation as well as with a fee-for-service scheme. The author finds that fee-for-service dramatically increases overtreatment rates compared to fixed payments and capitation, and this result holds even when fee-for-service is accompanied by a performance target. This is hardly surprising given that overtreatment does not lead to worse health outcomes for a patient; it only reduces the patient's income to the benefit of the physician.

Overall, a consistent pattern in this strand of the literature is that fee-for-service payment schemes lead to substantially higher rates of overtreatment compared to capitation and to fixed salaries. On the other hand, capitation and fixed salaries also cause an efficiency loss due to higher rates of undertreatment ([Hennig-Schmidt et al., 2011](#)). So the tradeoff involved in the comparison between those payment schemes are the same as those found in the comparison of integration vs. separation. This insight is also confirmed in a recent lab experiment by [Greiner et al. \(2017\)](#), which is framed as a patient-physician interaction. The authors compare outcomes between a baseline treatment,

¹⁰ The intuition for the first-best result is as follows: Without discounting, the consumer does not care whether her problem is solved in period 1 or period 2. Thus, a truthfully reported failure of the cheap treatment in period 1 causes real resource costs only for the expert who is obliged to solve the problem in period 2. This means that a fixed price contract (one where the payment from the consumer to the expert is independent of the treatment provided and independent of the customer's report) solves all the problems: The consumer has no incentive for misreporting (because the payment does not depend on the report and because receiving the expensive treatment in period 2 provides no additional utility if the cheap treatment provided in period 1 was sufficient) and the expert cannot profit from refraining from diagnosis or from untruthful reporting either (because he is essentially the residual claimant and as such fully liable for any misbehavior).

in which the subject in the role of physician makes a treatment recommendation and receives the price for this treatment, and a treatment with separation where one physician provides a diagnosis and treatment recommendation, while a different physician provides the treatment. Separation of diagnosis and treatment leads to a reduction in prices and in overtreatment rates (by 26%), but on the other hand it also increases undertreatment rates by 19%. As a result of the latter effect, separation is associated with only modest and insignificant efficiency gains of 5%.

4.2. The role of insurance

Insurance coverage for consumers is likely to have a profound effect on service provision in markets for credence goods, since it directly affects the consumer's residual costs of treatment and, thereby, it can indirectly also affect the expert seller's incentives.

Sülzle and Wambach (2005) investigate the relationship between the co-insurance rate (i.e., the share of the costs borne by the consumer) and the level of fraud in the market theoretically and arrive at the somewhat counter-intuitive conclusion that a higher co-insurance rate may be associated with either a higher or a lower level of fraud. In their model the consumer's problem is either minor (requiring a cheap treatment) or major (requiring an expensive treatment). The authors assume that the treatment provided is not observable by consumers (implying that overtreatment is dominated by overcharging), and that the expert is liable to solve the customer's problem (implying that undertreatment is ruled out). Thus, the focus of this study is the overcharging dimension of fraud. Prices are exogenously given with the price of the major treatment exceeding the price of the minor one. As a consequence, experts have indeed material incentives to overcharge their customers. Customers can either accept a given recommendation or obtain a second opinion from another expert. The benchmark model without insurance arrangements has three equilibria. One equilibrium is in pure strategies. In this equilibrium no customer searches for a second opinion and all experts are completely dishonest (recommending a major treatment to all customers who need only the minor one). In this ('full-fraud') equilibrium customers accept the expensive recommendation for sure because they know that all other experts are dishonest, too, and experts are fully dishonest, because they know that the expensive recommendation is accepted for sure. The other two equilibria are in mixed strategies. In those equilibria the extent of overcharging is just sufficient to induce some customers to seek second opinions, and customers seek second opinions sufficiently vigorously to prevent experts from overcharging all the time. In one of the two mixed strategy equilibria (the 'low-fraud' equilibrium) the market level of fraud is relatively low and few customers seek second opinions. In the second mixed strategy equilibrium (the 'high-fraud' equilibrium), the market level of fraud is relatively high and many customers seek second opinions. In the low-fraud equilibrium customers on their first visit do not reject expensive recommendations all the time (but only with positive probability), because the current diagnosis is likely to be correct, while in the high-fraud equilibrium first-time-customers do not reject expensive recommendations all the time because the next expert is likely to overcharge too.

Sülzle and Wambach (2005) show that the impact of introducing coinsurance depends on the equilibrium that is played in the no-insurance benchmark. When the starting point is the full-fraud equilibrium, then introducing coinsurance has no impact on the performance of the market. When the starting point is the low-fraud equilibrium, then introducing coinsurance leads to less fraud in the market and a lower probability of customers searching for a second opinion, and when the starting point is the high-fraud equilibrium, then coinsurance leads to more fraud and

more search. The intuition is that an increase in the coinsurance rate increases *ceteris paribus* the customer's incentive to reject an expensive recommendation, as she has to pay a larger fraction of the bill. So, in the low-fraud equilibrium experts have to become even more honest to keep customers indifferent between acceptance and rejection, while in the high-fraud equilibrium they have to become even less honest to keep the customer indifferent.

Some recent work has provided experimental evidence on the role of insurance in credence goods markets. Huck et al. (2016b) were the first to examine this topic experimentally. They conduct a lab experiment framed as a physician-patient interaction and model insurance as an experimental condition in which the costs of treatment by the physician are borne by all patients in a session collectively and not only by the affected patient. They find that insurance leads to more overtreatment because physicians anticipate that patients are less concerned about the costs of their treatment. Balafoutas et al. (2020) set up a theoretical model to investigate the effects of diagnostic uncertainty and insurance coverage for service provision in credence goods markets, and then test the model's predictions in lab experiments.¹¹ With respect to insurance, the authors find that it affects efficiency through a number of channels. On the one hand, and in line with the model's predictions, insurance coverage increases mistreatment rates and lowers investments in diagnostic precision by expert sellers. On the other hand, insurance increases market entry rates by consumers since they do not need to care for the costs. Overall, this latter effect is not sufficient to revert the negative relationship between insurance and efficiency resulting from the less consumer-friendly market behavior of the expert seller.

Turning to field evidence, the audit study of Lu (2014) in China shows that insurance leads to higher patient expenditure: those patients who are insured and signal that they plan to buy the prescribed drugs for diabetes or hypertension from the doctor's hospital also receive prescriptions for more expensive drugs, on average. A more recent paper by Balafoutas et al. (2017) builds on the experimental design by Balafoutas et al. (2013) and shows that fraud by taxi drivers in a natural field experiment in Athens becomes more likely and more extensive when the driver knows that the passenger's expenses are covered by his or her employer. Fraud takes the form of overcharging, which is far more likely (37% versus 20%) to occur when passengers indicate that they will be reimbursed for the costs of the ride. The authors attribute the cost increase for the client to 'second-degree moral hazard': if the service provider (in that case the driver) expects that the customer is not concerned about minimizing costs because a third party (in that case the employer, in other cases an insurance company) bears it, he may be more inclined to suggest or prescribe a more expensive service.

Kerschbamer et al. (2016) document a similar effect as Balafoutas et al. (2017) in a natural field experiment in the market for computer repairs. In a control treatment, the average repair price is about €70, whereas the repair bill increases by more than 80% when the service provider is informed that insurance would reimburse the bill. The experimental design in Kerschbamer et al. (2016) allows decomposing the sources of this impressive difference, showing that it is mainly due to the overprovision of parts and overcharging of working time. A survey among repair shops shows that the higher bills are mainly ascribed to insured customers being less likely to be concerned about minimizing costs because a third party (the insurer) pays the bill.

Summing up the existing evidence, the findings from Huck et al. (2016b), Kerschbamer et al. (2016), Balafoutas et al. (2020)

¹¹ The model is described in Section 5.

indicate that a comprehensive examination of the role of insurance in markets for credence goods should consider not only the demand side (behavior of customers) but also the supply side (behavior and fraud by expert sellers) on such markets. On the supply side, insurance coverage seems to have a pronounced fraud-enhancing effect, mainly through an increase in the overtreatment and the overcharging rate. On the demand side, insured consumers are more willing to accept. Therefore the overall effect of insurance on efficiency is less than obvious – depending on the circumstances it might well be positive (see Table 2 in [Huck et al., 2016b](#), for instance).

4.3. Liability and verifiability

Liability is defined as ‘the necessity for a seller to provide a good of sufficient quality to meet the consumer’s needs’ and verifiability as ‘the necessity for a seller to charge for the quality provided’ (DK06). The former of these two conditions precludes undertreatment, while the latter precludes overcharging. As discussed in the introduction, in the model of DK06 liability and verifiability both yield efficient outcomes (independently of whether they come in isolation or are both present) provided some further conditions are satisfied. One of these conditions is the commitment assumption. This assumption has two parts. First, commitment on the expert side and secondly, commitment on the buyer side. Commitment on the expert side means that the expert seller has the power to pre-commit *ex ante* that he will provide a treatment *ex post* (even if the price he has posted for the treatment does not cover the cost); and commitment on the buyer side means that once a recommendation is made, the customer is committed to undergo a treatment (even if the price she has to pay is large than her valuation). If the commitment assumption is violated, then inefficiencies and fraud might arise in the market equilibrium even when liability or verifiability is in place. We have already discussed the inefficiencies that may arise if commitment on the buyer side is violated in the verifiability but no liability case (see the discussion of [Fong et al., 2014](#) in 3.2). As we have seen there, the equilibrium is particularly inefficient in those cases where the problem associated with the greater surplus occurs with a larger probability: When the surplus of solving the major problem is larger than the surplus of solving the minor problem and the major problem is sufficiently likely, the most profitable equilibrium involves full overtreatment; in the opposite case, where the minor problem is associated with the larger surplus and is very frequent, the most profitable equilibrium involves full undertreatment. In all other cases, the problems are either less severe or the equilibrium involves honest recommendations and efficient treatment.

[Fong et al. \(2014\)](#) compare these equilibrium outcomes (for the verifiability but no liability case) to those resulting from the same model when verifiability is replaced by liability. For the no-commitment case, the market with liability has previously been analyzed by [Fong \(2005\)](#) and yields the following outcomes: If the expected benefit from successfully solving the customer’s problem is larger than the cost of the expensive treatment, then the market is fully efficient: In this case the expert can extract the entire expected surplus from solving the problem by posting a uniform price equal to the expected surplus.¹² This is no longer feasible if the expected benefit from successfully solving the customer’s problem is smaller than the cost of the expensive treatment. The market equilibrium in this case is one

in mixed strategies, in which the expert posts two prices and recommends honestly because the minor recommendation is always accepted while the major recommendation is sometimes rejected.¹³ In this equilibrium, inefficiencies arise because the customer has to reject expensive recommendations sufficiently often to remove the expert’s incentives for overcharging. [Fong et al. \(2014\)](#) compare the inefficiencies across the two environments for different parameter constellations. Their main finding is that under the liability assumption the market outcome is more efficient than under the verifiability assumption when the problem associated with the greater surplus is sufficiently more likely, while the opposite relationship might hold for other parameter constellations.

In a large-scale lab experiment which has since become paradigmatic for the study of credence goods markets in the economic laboratory, [Dulleck et al. \(2011\)](#) investigate how informational and institutional restrictions (like liability or verifiability) and market conditions (like competition and the possibility to build up a reputation) affect behavior on a credence goods market. In the experimental market studied by [Dulleck et al. \(2011\)](#), sellers and buyers interact over 16 rounds in a typical credence goods market where the seller has superior information about the kind of problem faced by the consumer, and hence about the kind of solution that best fits this problem. Consumers decide whether to trade on the market (and once they enter the market, they must accept the seller’s recommended treatment following the commitment assumption), while sellers decide on which treatment to provide and on which treatment to put on the bill for each consumer they serve on the market.

The authors find that liability has a crucial, but verifiability at best a minor effect. Indeed, while the relative frequencies of market interaction, undertreatment and overtreatment do not differ significantly between the baseline treatment without verifiability and liability and the verifiability treatment, the market performs significantly better under liability. Also, the observed aggregate performance in the two treatments without liability is better in terms of efficiency than the standard prediction for a market without verifiability and liability, but considerably worse than the prediction for a market with verifiability. The findings in [Dulleck et al. \(2011\)](#) raise at least three questions: Why is the performance of credence goods markets so poor in the presence of verifiability where theory predicts full efficiency? Why does the market without verifiability and liability perform so much better than predicted? And why are the results for the market with liability well in line with the theoretical prediction?

Those questions are addressed in [Kerschbamer et al. \(2017\)](#). In the theory part of the article, the authors argue that heterogeneity in the social preferences of credence goods sellers can provide an answer to all of those questions. Key to their argument are the following observations: First, the standard solution to the credence goods problem for markets with verifiability but without liability – equal-mark-up prices – is robust to the presence of sellers with pro-social other-regarding preferences, but non-robust to the presence of sellers with anti-social other-regarding concerns. Second, for the prediction for markets without verifiability and liability the opposite is true: It is robust to the presence of sellers with anti-social other-regarding preferences, but non-robust to the presence of sellers with pro-social other-regarding concerns. Third, liability is a robust institution – it performs as predicted independently of whether the seller is selfish or motivated by pro- or anti-social preferences.

A key ingredient in this explanation is heterogeneity in social preferences in the (experimental) seller population. To provide

¹² With a uniform price there is no overcharging and, since the customer receives no information from the expert about the nature of the problem, she accepts the treatment with probability one because the expert is liable and the price is not higher than the expected benefit from treatment.

¹³ The mixed strategy equilibrium is very similar to the one characterized by [Sülzle and Wambach \(2005\)](#) and discussed in Section 4.2.

support for heterogeneity, the authors design a test that allows for a clean discrimination between different preference types, based on the provision behavior in an experimental market for credence goods. The results obtained in an implementation of the test indicate that less than a quarter of the experimental sellers behave in accordance with the standard assumption on preferences. The rest behave either in line with other forms of selfish or in accordance with different variants of non-selfish social preferences. Taken together, these results provide strong support for heterogeneity in social preferences and therewith for the proposed explanation for the failure of verifiability to increase efficiency.

An alternative explanation for the failure of verifiability to yield efficiency is provided by Hilger (2016). In Hilger's model the consumer's problem is again either minor or major, the treatment provided is observable by consumers and the expert is not liable to solve the customer's problem. So far the model is identical to the one in DK06. The key innovation in Hilger (2016) is the assumption that consumers do not observe experts' cost functions. Specifically, there are two types of experts, minor specialists who have a comparative advance in solving the minor problem and major specialists who have an advantage in solving the major problem. Consumers know the relative frequency of the two types but cannot distinguish them. Hilger shows that under this assumption there exists no pure strategy equilibrium involving full honesty. The intuition is as follows. Since the two types of experts have different cost structures and since efficient provision requires equal mark-ups, full honesty requires separation (that is, in a fully honest equilibrium the two types of expert must post different price vectors, each implying equal mark-ups for the respective cost structure). But under separation it is necessarily the case that one of the two types of expert has an incentive to mimic the other type. That type then either has an incentive to always recommend the minor or to always recommend the major treatment.

Summing up this subsection, it seems fair to say that, although verifiability and liability both yield full efficiency under the specific assumptions of DK06, liability outperforms verifiability in terms of robustness – in the sense that it increases efficiency under a broader set of alternative assumptions. It should be added, though, that there are many limits to the implementation of liability rules in practice. This is especially true in the medical realm. As discussed by DK06, for many diseases there is no sufficient treatment, while with others success is in part random. Thus, a failing treatment is no perfect signal of undertreatment. Also, liability requires that the success of a service is not only observable by the patient but also verifiable. This is often unrealistic, especially in the medical realm. Strict liability rules have their limits also outside the medical sector. For example, a car might work when it is taken up from the mechanic, but it might show the same symptoms after a few weeks. To prevent undertreatment in such a case, liability would have to cover a longer time period. But during the longer period the car might break down for other reasons. Also, strict liability rules that cover a longer time period might invite moral hazard on the consumer side. For all those reasons, liability might work better in theory than in practice.

4.4. Competition between sellers

Competition between sellers, either in the form of price competition or competition for buyers through the provision behavior, appears as a promising feature of the institutional setup that could restrain fraud by expert sellers and lead to higher efficiency. Inderst and Ottaviani (2012c) build competition into their model of financial advice as a credence good, in which informed financial advisors interact with consumers who are either

naïve (anticipating that advisors offer unbiased advice) or wary (anticipating that advisors are driven by a desire to maximize commissions). They show that, when two advisors compete for consumers, consumer surplus and efficiency increase. This is the case both when consumers are wary and when they are naïve.

In the experiment of Dulleck et al. (2011), competition between sellers is found to have an impact on efficiency when combined with liability but not if liability is absent. Competition in this experiment is operationalized as an option for consumers to choose among several sellers, as opposed to exogenous bilateral consumer–seller matches. The experimental results are largely in line with the theoretical predictions, which can be summarized as stating that introducing competition increases the frequency of trade (compared to the baseline) and leads to lower prices, but it does not guard consumers against undertreatment. Indeed, in the experiment, the frequency of trade increases substantially as a result of adding competition. This is true in all four liability – verifiability constellations. Moreover, prices drop when sellers compete for buyers, but undertreatment does not decrease (and it even goes up from 53% to 73% when neither liability nor verifiability are in place). Hence, overall, competition drives down prices and increases trade. However, it does not lead to higher efficiency as long as liability is violated, because sellers frequently provide an insufficient service in this case.

In a recent study using field data from undercover car repair shop visits carried out by the German Automobile Association (ADAC), Rasch and Waibel (2018) examine the effect of competition between sellers on fraud. Competition is operationalized as market concentration, i.e., the number of car repair shops within 10 kilometers of a shop. The authors report that a high degree of competition is associated with a significantly lower likelihood of fraud.¹⁴ This finding is in line with the findings of Huck et al. (2012, 2016a). These studies use lab experiments to investigate the impact of competition in markets for experience goods (as compared to credence goods) and find that the possibility to choose between sellers reduces the occurrence of fraud.¹⁵

Mimra et al. (2016a) combine elements from Dulleck et al. (2011) and Huck et al. (2012, 2016a) in their investigation into how the intensity of price competition and the level of consumer information influence market efficiency when experts can build up reputation. With respect to competition, a rather surprising result emerges: the authors find that the level of fraud in the form of undertreatment and overcharging is higher under price competition than under fixed prices: In the early periods (in which reputational concerns may play a role), customers return significantly less often to experts who have undertreated them under fixed prices than under price competition. Thus, for given prices punishment is higher in the fixed price condition. Furthermore, under price competition prices are on average lower than in the fixed price condition and this price pressure undermines reputation-building further because the benefits of a good reputation are lower when the prices are lower.

Based on all of the above, we sum up this subsection by saying that the impact of competition on fraud in credence goods markets is at best ambiguous.

¹⁴ A similar measure of competition as in Rasch and Waibel (2018) is used in Gottschalk et al. (2020), but the results are different: Gottschalk et al. (2020) report that dentist density (the number of dentists within a 500 meters radius) does not affect overtreatment in their experiment.

¹⁵ It should be noted here that experience goods differ from credence goods in several important dimensions. For example, (1) while the valuation of a consumer is strictly increasing in quality with experience goods, it is constant whenever the quality is sufficient with credence goods; (2) for given prices a consumer can tell exactly which quality he prefers in the case of experience goods, but he does not know it with credence goods; (3) whereas the quality of the good is unobservable ex ante but perfectly observable ex post with experience goods, it may be observable either ex ante, or ex post, or neither ex ante nor ex post with credence goods.

4.5. Reputation

The possibility that fraud by expert sellers may be restrained if they are concerned about their reputation has been discussed frequently in the literature. For instance, the seminal work by [Kreps and Wilson \(1982\)](#) suggests that reputational concerns by expert sellers might lead to more consumer-friendly behavior on markets for credence goods. Their idea is that an expert seller who knows that his performance is observed by potential future customers might refrain from exercising fraud because it is observable in the long run. Indeed, it is quite intuitive to think that revealing information about the past behavior of experts can discourage them from defrauding their consumers in order to secure future interactions with these – and other – consumers. A theoretical study along these lines is [Frankel and Schwarz \(2014\)](#) who show that sometimes – although not always – experts whose past actions are publicly observable will be led to choose the appropriate treatment for the consumer.

In their large-scale lab experiment on the influence of institutional and market conditions on the extend of fraud in credence goods markets, [Dulleck et al. \(2011\)](#) also have a (private-history) treatment, in which they allow consumers to keep track of their own past experience with a particular seller. This creates the possibility for reputation building on the part of expert sellers. From the theoretical model underlying their experiments, [Dulleck et al. \(2011\)](#) derive the prediction that the possibility for reputation building does not affect market outcomes (interaction decisions by consumers and provision and charging decisions by sellers) unless competition between experts is also part of the setting. If consumers can reward honest experts by selecting them as interaction partners, then the possibility for reputation building is predicted to increase efficient trade. The authors report that adding the possibility for reputation building indeed increases efficiency thanks to a higher volume of trade, but only when neither liability nor verifiability are in place. Moreover, under the same conditions, reputation reduces the frequency of overcharging.

In a related lab experiment based on the design of [Dulleck et al. \(2011\)](#), [Mimra et al. \(2016a\)](#) find that public information about the past behavior of expert sellers does not necessarily reduce the level of fraud. Similarly, [Huck et al. \(2016a\)](#) show that replacing private by public histories has only a minor effect in an experimental market for experience goods.

Turning to experimental evidence from the field, [Schneider \(2012\)](#) investigates in his pioneering study the impact of reputational concerns on market performance. Reputation in [Schneider \(2012\)](#) is operationalized through the undercover consumer sending a signal to the car mechanic that future interaction is likely (by mentioning that the consumer is new in town and looking for a reliable garage, as opposed to mentioning that she is leaving town). Based on data from undercover garage visits the author finds that overtreatment and undertreatment in the car repair sector are pervasive but he finds no evidence indicating that reputation considerations affect mechanics' provision or charging behavior. The only evidence the author finds for the impact of reputational concerns is a change in the average diagnosis fee – which is approximately 37% lower when reputation matters. A plausible interpretation put forward by the author is that car mechanics choose to respond to reputation concerns along the fee dimension rather than along the provision dimension, because the latter is more easily and directly observable by customers.

In addition to the above studies that attempt to answer the question of whether reputation and the prospect of future interaction can have a fraud-reducing effect on credence goods markets, part of the literature considers the possibility that reputational concerns might also lead to less and not only to more consumer-friendly behavior (see [Ely and Välimäki, 2003](#); [Ely](#)

[et al., 2008](#) for formal models).¹⁶ [Grosskopf and Sarin \(2010\)](#) explicitly test for the presence of 'good' as well as 'bad' reputation in a series of lab experiments. They employ a clever design with two types of players interacting with each other: long-run players represent sellers whose past actions may be observable depending on the treatment, while short-run players represent consumers who live only for one period. Experts belong to one of two types, either 'friendly' or 'unfriendly'. The main finding in [Grosskopf and Sarin \(2010\)](#) is that the positive effects of reputation on market efficiency are generally not as strong as predicted by theory, while they find hardly any evidence of a negative effect.

Taken together, the studies discussed in this subsection fail to find a significant or particularly systematic impact of reputation and information revelation on fraud in credence goods markets.

4.6. Second opinions

The idea behind promoting the search for multiple opinions as a potential remedy to the credence goods problem resembles to some extent the discussion of the potentially fraud-reducing effects of competition between experts and separation of diagnosis and treatment. The intuition is straightforward: If the expert seller who performs the diagnosis is not a monopolist and consumers have the possibility of seeking more (or better) advice from another expert, then – depending on the level of search costs – the pure threat of second opinions in the market might be a valid instrument to reduce fraud. A relevant recent example is legislation in Germany, which since 2016 allows patients to receive a costless second opinion for certain surgical procedures that have been shown to be particularly prone to abusive practice by physicians, such as spinal column operations.¹⁷ In Switzerland, some insurance companies even offer discounts of 10 to 15 percent to customers who commit to seek a second opinion before they undergo certain types of surgery.¹⁸

The pioneering theoretical work studying the role of gathering multiple opinions in mitigating fraud in markets for credence goods is [Pesendorfer and Wolinsky \(2003\)](#). The problem on which the authors focus is not mistreatment or overcharging (as in much of the literature discussed thus far) but rather moral hazard in the provision of diagnosis effort. Specifically, the authors consider a market in which the expert must exert costly effort to identify the service that meets the consumer's need, the expert's diagnostic effort is unobservable and the final success of the service is not contractible. The set of possible services is modeled as a continuum to ensure that an unguided guess will not yield the appropriate service with positive probability. All types of the service are equally costly and all have the same price so that experts have no incentive to misrepresent the correct diagnosis in case they know it. Experts are sampled sequentially from a large population and sampling an expert is costly for the consumer. A sampled expert offers the consumer a contract consisting of a price for the diagnosis and a price for the service. Upon observing the contract, the consumer decides whether to receive a diagnosis from the expert or to continue sampling. Once the consumer agrees to be diagnosed, the expert decides whether to incur the cost of effort and then provides a recommendation. After learning the recommendation – but without knowing whether the expert

¹⁶ Consider, for instance, a car mechanic who – by chance – had to replace the engine in the previous two repairs. If for the next car the appropriate repair is again replacing the engine, the mechanic might refrain from the appropriate repair in order to avoid getting the reputation as a mechanic who always replaces the engine.

¹⁷ See https://dejure.org/gesetze/SGB_V/27b.html.

¹⁸ See <https://www.comparis.ch/krankenkassen/info/glossar/zweitmeinung-second-opinion.aspx>.

has exerted effort – the consumer either buys the service or continues her search.

The authors show that equilibria with positive diagnosis effort necessarily involve mixed strategies. The reason is that if all experts invest in accurate diagnosis, then the consumer has no incentive to incur the cost of acquiring a second opinion. But if the consumer never acquires a second opinion, experts have no incentive to invest in diagnosis. Specifically, the authors show that if prices are exogenously given (and the search cost and the effort cost are not too large) then the model exhibits one pure strategy equilibrium and two mixed strategy equilibria – similar to the model by [Sülzle and Wambach \(2005\)](#) discussed in [4.2](#): The pure strategy equilibrium is a ‘full fraud’ equilibrium in which experts do not invest in diagnosis and the consumer quits immediately. The mixed strategy equilibria are ‘partial fraud’ equilibria (one with a higher level of fraud than the other) in which the experts choose high diagnostic effort with positive probability, since high diagnosis effort ensures that the consumer will accept the recommendation with positive probability, while with low effort the consumer rejects for sure. This latter property is due to the fact that the consumer searches for two matching recommendations, and with infinitely many possible services the recommendations of two experts can only match if both experts have exerted effort.

[Agarwal et al. \(2019\)](#) test (in a simplified framework) some of the core predictions of the [Pesendorfer and Wolinsky \(2003\)](#) model in lab experiments. They show that – compared to a baseline market with no institutions for information acquisition prior to trade – introducing the possibility to search for a second opinion significantly increases the likelihood of trade, the probability of correct treatment and market efficiency. The authors also examine the effect of a change in the search cost of the consumer on market efficiency. Consistent with theoretical predictions, lowering the search cost yields a significantly positive effect on market efficiency.

A closely related experimental test of the role of second opinions and how these interact with the search costs for consumers is conducted by [Mimra et al. \(2016b\)](#). However, while the experiments in [Agarwal et al. \(2019\)](#) are based on the [Pesendorfer and Wolinsky \(2003\)](#) model, the experiments in [Mimra et al. \(2016b\)](#) are based on the [Sülzle and Wambach \(2005\)](#) model. Remember that a core assumption in [Sülzle and Wambach \(2005\)](#) is that the service provided is not observable by consumers and that the expert is liable to solve the customer’s problem. Thus, this model’s focus is on overcharging. [Mimra et al. \(2016b\)](#) change this assumption by looking at markets where there is liability and verifiability and where prices are exogenously fixed such that selling the major treatment is more profitable than selling the minor one. Thus, their focus is on the overtreatment dimension of fraud. The rest is exactly as in [Sülzle and Wambach \(2005\)](#): Again there are three equilibria, one ‘full fraud’ equilibrium in pure strategies and two ‘partial fraud’ equilibria in mixed strategies.

The authors find that introducing the possibility of searching for a second opinion significantly reduces the level of overtreatment (compared to a baseline where this option is absent). Market efficiency increases significantly with the introduction of second opinions when search costs are low, but not when search costs are high. The latter effect is due to the fact that there are two opposing effects on market efficiency: The reduction in overtreatment (which increases efficiency through the reduction in treatment costs) must be weighed against the search costs incurred by consumers whenever they elicit a second opinion.¹⁹

¹⁹ In a different context, [Dulleck and Wigger \(2015\)](#) show in their model of the political market that politicians as experts may be restricted in the abuse of their informational advantage over the electorate through the presence of political opposition. Hence, in this model the opposition in effect plays the role of an expert who can provide a second opinion regarding the true state of the economy, leading to higher welfare for voters.

In their field experiment on the Indian market for life insurance discussed earlier, [Anagol et al. \(2017\)](#) also run a treatment in which the confederates stated to the life insurance agent that they have previously received advice from another agent who recommended a particular type of insurance. The authors find that, overall, mentioning that one has received advice from another agent does not affect the quality of advice by agents in the experiment. However, when this second opinion was a poor recommendation, there is an increase in the quality of advice. Hence, there is some evidence indicating that mentioning a second advisor improves the quality of advice one gets on the market. Nevertheless, it not clear whether this positive effect is due to competition between advisors (meaning that some advisors de-bias consumers as a way of demonstrating their superior competence compared to their competitors), or simply due to advisors trying a different recommendation than the one that their customer has already rejected from another agent.

In their field study on dental care, [Gottschalk et al. \(2020\)](#) have a treatment variation that allows them to examine how (the threat of) second opinions affects fraud by professional dentists. In one condition, the patient indicates to the dentist that he has uploaded his dental X-ray to an internet platform where dentists offer free advice, and that he is awaiting a response. Hence, the patient in this condition signals that he is likely to receive a second opinion from another expert in the near future. The authors find no treatment difference in this respect, with overtreatment rates decreasing only slightly and insignificantly, from 29% to 27% in the second opinion condition.

That second opinions may – at least in some contexts – reduce fraud and increase consumer welfare is only one side of the story: [Alger and Salanie \(2006\)](#) show that giving the consumer the option to search for a second opinion might give rise to a perverse equilibrium in which expert sellers always recommend an extensive treatment in order to keep customers uninformed, thereby preventing them from seeking a better price elsewhere. A similar result is obtained by [Dulleck and Kerschbamer \(2009\)](#) in a different context: They characterize an equilibrium in which experts undertreat consumers with strictly positive probability to keep them uninformed as this deters them from free-riding on experts’ advice. Although these results sound very similar, the forces leading up to them are different.

Both papers study a simple credence goods market in which consumers have a loss that requires either a minor or a major treatment. In [Alger and Salanie \(2006\)](#) the quality of the service provided by the expert is only partially verifiable.²⁰ If the quality is not verifiable, experts have to post the same price for both qualities to avoid fraud, and if verifiability involves a cost the prices have to be similar. Since the costs of the two qualities differ but the price is the same or similar, there is a cross-subsidization from low to high quality. If the cost of getting a second advice is low, this cross-subsidization invites cream skimming by a deviating expert aimed at consumers who have learned that they need the minor treatment. In equilibrium this problem is solved by having the expert always recommend the same intervention so that the diagnosis is uninformative. This solution is costless (and not very interesting) if the quality of service is not verifiable (the expert simply has to lie to the consumer), but it becomes interesting if the provided treatment is somewhat verifiable so that pooling entails a (fraud) cost. This is the setting [Alger and Salanie \(2006\)](#) focus on.

By contrast, [Dulleck and Kerschbamer \(2009\)](#) consider a market where the quality of service is verifiable, but where costly

²⁰ Specifically, the authors allow the degree of verifiability of the quality to change continuously; the result we are discussing here is obtained for a low degree of verifiability.

(and unobservable) diagnosis effort is needed to identify the consumer's needs. With unobservable diagnosis, the diagnosis cost has to be carried by the price for the low quality service to avoid fraud, again implying a cross-subsidization. Like in the Alger and Salanié model, this cross subsidization induces experts to lie to consumers to keep them uninformed, thereby preventing them from seeking a better price elsewhere.²¹

Summing up the discussion in this subsection we conclude that the impact of allowing for multiple opinions on the level of fraud in credence goods markets is ambiguous: There are forces that lead to less fraud and higher efficiency and countervailing forces that lead to more fraud and less efficiency.

5. Additional dimensions of the seller–consumer interaction: Diagnostic effort, prices, and tax attitudes

Most of the literature on credence goods markets has assumed that an expert perfectly observes the consumer's problem and can subsequently recommend a treatment to solve this problem. This amounts to assuming that diagnosis is not only perfectly accurate, but also – in the large majority of studies – that it is costless to the expert. Given that in practice diagnoses are often imprecise and require a certain amount of effort invested by the expert (with healthcare services being a prime example of such a situation), it is interesting to consider diagnostic effort – and the resulting diagnostic precision – as an additional variable of interest.

The first study to theoretically integrate diagnostic effort into the setting of expert services was [Pesendorfer and Wolinsky \(2003\)](#). Since we have already discussed this paper in some detail in Section 4.5, we do not discuss it further here. [Dulleck and Kerschbamer \(2009\)](#) – also shortly described in the previous subsection – consider a standard credence goods market in which the consumer's problem is either minor (requiring a cheap treatment) or major (requiring an expensive treatment). Regarding verifiability and liability the authors assume that the treatment provided is verifiable and that the expert is not liable to solve the customer's problem. An important feature of their model is that identifying the consumer's problem requires a costly diagnosis and that the expert's diagnostic effort is not observable by the consumer. Another important feature is that experts face competition by discounters who offer the two treatments at marginal cost without providing advice. The authors first derive the first-best benchmark in which all information is symmetric. Depending on the probability of having the major problem and on the diagnosis cost the efficient way to serve a customer is one of the following: Policy A: Performing a costly diagnosis and providing the quality matching the needs of the customer. Policy B: Performing no diagnosis and blindly providing the minor treatment. Or Policy C: Performing no diagnosis and blindly providing the major treatment. Intuitively, Policy A is the efficient solution if the diagnosis cost is sufficiently low and if the probability of the major problem is neither close to zero nor close to one.

With asymmetric information, three incentive problems burden the relation between consumers and experts: (i) A moral hazard problem on the expert side: how can the expert be induced to perform the costly but unobservable diagnosis? (ii) A truthful reporting problem on the expert side: how can the expert be induced to reveal the diagnosis outcome truthfully? And (iii) A free-rider problem on the consumer side: How can the consumer be prevented from taking the expert's advice and buying the

recommended treatment from a discounter? Under verifiability, solving problem (ii) requires equal mark-up prices; but with equal mark-ups (and no other forces pushing towards honesty) experts have no incentives to invest in diagnosis. To give experts a fair chance to solve both problems (i) and (ii) simultaneously, the authors give them the option to offer a warranty payment for the case where the cheap treatment fails.

[Dulleck and Kerschbamer \(2009\)](#) show that the efficient solution is sustainable in equilibrium if and only if the search cost of the consumer is sufficiently high compared to the diagnosis cost. To understand this result, suppose that the parameters of the model are such that Policy A is the efficient solution. To induce the expert to choose this policy, it must yield a higher profit than policies B and C. The temptation to choose Policy B is easily eliminated by committing to a high enough warranty payment for treatment failure. Policy C is more critical: To remove the temptation to opt for it, experts have to post prices that finance the diagnosis cost through the mark-up on the minor treatment. But this makes them vulnerable to competition by discounters. In a pure strategy equilibrium the consequence of the resulting free-rider problem is that (for some parameter constellations) consumers are served by discounters, although efficiency would require performing a costly diagnosis and providing the quality matching the needs of the customer. The authors also characterize a mixed strategy equilibrium in which experts cheat to keep consumers uninformed and in which the less precise information prevents consumer from seeking a better price elsewhere (this is the equilibrium we have discussed in the previous section).

[Bonroy et al. \(2013\)](#) extend the model of [Dulleck and Kerschbamer \(2009\)](#) by considering risk-averse consumers. Their main finding is that when consumers are risk averse, then even in the absence of competition by discounters, the market might not end up with Policy A when that policy is efficient. The reason is that the mark-ups needed to provide incentives for the expert to choose Policy A impose a (risk) cost on risk averse consumers. Thus, the region where Policy A is provided shrinks – in comparison to the first-best benchmark. The authors also show that competition between experts reduces the inefficiency if consumers' risk preferences satisfy decreasing absolute risk aversion. The reason is that competition drives mark-ups down; and that with lower mark-ups the range of parameters for which Policy A is efficient but not provided shrinks.

One further interesting recent study that has theoretically modeled costly diagnostic effort is [Bester and Dahm \(2018\)](#). We have already discussed some important features of that study in Section 4.1. Here we add some further details. As in [Dulleck and Kerschbamer \(2009\)](#), identifying the consumer's problem requires costly effort and the diagnostic effort undertaken by the expert is not observable to the consumer. The most important difference to [Dulleck and Kerschbamer](#) regards the verifiability of success or failure of the cheap treatment: In [Dulleck and Kerschbamer](#) success or failure of the low-cost treatment is assumed to be observable and verifiable; this allows for tariffs where the payment for the cheap treatment is contingent on success or failure. By contrast, in [Bester and Dahm \(2018\)](#) success or failure of the low-cost treatment is assumed to be observable but not verifiable; this allows for tariffs where the payment for the cheap treatment is contingent on the consumer's report about success or failure. This latter difference implies that, instead of the free-rider problem on the consumer side considered by [Dulleck and Kerschbamer \(2009\)](#), [Bester and Dahm](#) have a subjective-evaluation problem on the consumer side: How can the consumer be induced to report success or failure truthfully?

The authors show that, for some range of parameter combinations, the moral-hazard problem on the expert side, the

²¹ In the model by Alger and Salanié lying occurs in a pure strategy equilibrium and involves experts always claiming that the high quality is needed regardless of the customer's true needs (i.e., lying involves overtreatment). In the [Dulleck and Kerschbamer](#) model lying occurs in a mixed strategy equilibrium and involves experts sometimes claiming that the low quality is needed regardless of the customer's true needs (i.e., lying involves undertreatment).

truthful-reporting problem on the expert side and the subjective-evaluation problem on the consumer side can be solved simultaneously with equal mark-ups on the expected treatment costs before receiving information. The range for which the first-best can be implemented with equal mark-ups increases when the discount factor increases and includes all parameter combinations when the discount factor is one.

In the context of financial advice, [Inderst and Ottaviani \(2012c\)](#) consider an extension of their model where the advisor must invest costly effort in order to improve the quality of his privately observed information. They show how the optimal amount of invested effort depends on the advisor's prior beliefs and on commissions.

In a recent paper already mentioned in Section 4.2, [Balafoutas et al. \(2020\)](#) introduce diagnostic uncertainty into the model by DK06. In particular, they assume that the expert receives an imperfect signal about the actual problem of the consumer and has the possibility to reduce diagnostic uncertainty through costly investments in diagnostic precision. Another important feature of this model is that it includes insurance coverage for the consumer, who has to pay a fair insurance premium in exchange for full coverage of her expenditure on the market. Finally, the model allows for altruistic experts who potentially care positively for the material payoff of their customers. This model provides a framework for testing the effects of diagnostic uncertainty and insurance coverage on service provision and fraud by expert sellers and on market efficiency. The model predicts that experts are less likely to follow the signal they receive about the appropriate solution (i.e., they are more likely to engage in fraud) when diagnostic precision is relatively low and when insurance is in place. In addition, there is an interaction between the two dimensions, with the model predicting that exogenously introducing insurance into the setting lowers the investment of experts into diagnostic precision.

[Balafoutas et al. \(2020\)](#) test these predictions in a laboratory experiment that builds on the design by [Dulleck et al. \(2011\)](#), thus providing the first empirical evidence on the role of diagnostic uncertainty for credence goods provision. The experimental findings confirm the key predictions of the model: Diagnostic uncertainty not only increases fraudulent provision rates, but it also reduces market entry rates by consumers. These effects combined have a strong negative impact on efficiency, reducing it by about 16%.

Another dimension that is worth discussing is the choice of posted prices by expert sellers. [Dulleck et al. \(2012\)](#) jointly analyze posted prices and cheating behavior by experts, with the purpose of addressing the question of causality between pricing and cheating behavior. Starting point of this study is the observation that in the lab experiments by [Dulleck et al. \(2011\)](#) a non-negligible fraction of experts 'do the right thing' even when material incentives are strongly distorted. For instance, for an environment without institutional safeguards against (and therefore strong material incentives for) fraud, [Dulleck et al. \(2011\)](#) find that almost 30% of experimental experts provide the appropriate quality throughout the 16 rounds of the experiment. They also present evidence showing that 'fair prices' – that is, prices that split the gains of trade equally between the two parties of a transaction – are much more frequently observed than predicted, and that they are atypically often associated with efficient provision behavior despite the fact that they provide material incentives for undertreatment. Those observations raise the question: do 'fair prices' induce 'good behavior', or do 'good experts' post 'fair prices'? The answer to this question has important practical implications, for instance, for regulation policy: If fair prices induce good behavior, then price regulation might be a cheap and effective means to increase efficiency in markets

for credence goods; however, if good experts post fair prices, then price regulation may actually reduce efficiency, as the prices posted by expert sellers in unregulated markets contain valuable information for customers about the type of seller they face, information that would be destroyed by regulation.

[Dulleck et al. \(2012\)](#) address the causality issue by first presenting a model with heterogeneous sellers, and then running new experiments to test the main implications of the model. The theoretical model implies a causal effect from other-regarding preferences towards pricing behavior and not vice versa. The causal unidirectionality is confirmed empirically, by comparing experimental conditions in which prices are endogenous, i.e. chosen by the expert sellers, to conditions where prices are exogenous, i.e. randomly set by the experimental software. The presented evidence suggest that, with endogenous prices, whether a transaction takes place under a fair or an unfair price-vector significantly affects both the probability with which the consumer accepts and the probability that she receives efficient service. However, those effects disappear when controlling for expert type with fixed effects, or by exogenously imposing prices. Thus, it is not the case that fair prices induce good behavior; instead, good experts frequently choose fair prices.

Posted prices are also part of the analysis in the paper by [Mimra et al. \(2016a\)](#) on the effects of competition and reputation on the level of fraud in a credence goods market. Unsurprisingly, price competition in that experiment leads to lower prices compared to a non-competitive benchmark with exogenously given prices. Interestingly, undertreatment and overtreatment are significantly higher with endogenous prices than when prices are exogenously fixed.

[Balafoutas et al. \(2015\)](#) consider the interaction between fraud (in the form of undertreatment and overcharging) and revealed tax attitudes in credence goods markets. The motivation behind [Balafoutas et al. \(2015\)](#) is that tax evasion often requires the mutual consent of a seller and a buyer of a good or service, which has the consequence that both agents have the opportunity to find out each other's attitudes towards tax evasion before the actual product or service is traded. This can in turn influence market outcomes. For instance, a buyer matched with a seller who has revealed that he would be willing to hide the transaction from the tax authorities might interpret this information as a signal of the seller's general dishonesty and abstain from interacting with him. The experimental results strongly confirm that there is a link between revealed tax attitudes and market outcomes: when at least one of the contracting parties (the buyer or the seller) is tax-dishonest, efficiency is substantially lower compared to the situation in which two tax-honest agents interact.

6. Concluding remarks

The literature summarized in this review takes as the defining feature of a credence good the asymmetric information between the expert seller and his customer regarding the fit between the characteristic of the product and the needs of the customer. There is a second strand of literature that defines the term credence good differently – namely as a good or service that has unobservable attributes that remain undetected even after consumption. In this alternative definition the asymmetric information concerns the characteristics of the good, and not its suitability for a particular consumer. To distinguish goods defined in this way from classical credence goods, [Kerschbamer \(2018\)](#) call them *label credence goods*. Important examples for label credence goods are organically produced food, fairly traded products, 'dolphin-safe' caught tuna, or products containing no genetically modified ingredients.

Label credence goods are also widespread, and their importance seems to have increased over time (see, for instance, the

increasing discussion about Corporate Social Responsibility in the media). Despite their increasing importance, we have chosen to not include label credence goods in our review, since the main research questions covered in this literature and the proposed remedies against fraud are different. For instance, for standard credence goods markets where the success of service is observable and verifiable, imposing liability rules is a potentially powerful remedy against fraud. Since the 'success of service' is almost by definition not observable in a label credence goods market, this solution is not discussed in this second strand of literature. Also, while the search for multiple opinions might be a powerful means to contain fraud in a standard credence goods market where the asymmetric information regards the fit between the properties of the good and the needs of the customer, this remedy has no obvious counterpart in a label credence goods market where the customer typically knows what he would like to consume but cannot assess whether a given good has the desired characteristics.

Turning to the remedies against fraud, important forces affecting sellers' behavior discussed in the label-credence-goods literature are self-labeling through unverifiable quality claims (see, e.g., Cason and Gangadharan, 2002; Baksi and Bose, 2007, or Etilé and Teyssier, 2016), disclosure of verifiable quality information (c.f., e.g. Dranove and Jin, 2010 or Brenndorf et al., 2015), quality labels certified by third parties (see, e.g., Cason and Gangadharan, 2002; Baksi and Bose, 2007; Dranove and Jin, 2010; Harbaugh et al., 2011; Farhi et al., 2013), minimum quality standards imposed by the government (Glaeser and Ujhelyi, 2010), voluntary quality-enforcing organizations formed by firms (Baron, 2011; Fischer and Lyon, 2014) and information-supplying activists (Feddersen and Gilligan, 2001; Lyon and Maxwell, 2011).

In more general terms, we believe that the potential for fraudulent behavior by experts who possess superior information is present in many markets beyond those traditionally referred to as examples of classical credence goods (such as healthcare or repair services) or label credence goods (such as green ingredients in food or special production standards in all kinds of goods). We have already outlined in our review the new exciting developments in the field of finance, which often analyze the provision of financial services in a classical credence goods framework. There are many services provided in this sector that would perfectly fit in the label-credence-good category as well (think of the emerging field of ethical investment, for instance). The same is true for environmental goods and services. Additional fields could include effort or information provision by politicians (see also Section 2 of this review), journalists, or even religious leaders or representatives. In modern economies that are characterized by the rapid development of new kinds of services, by highly specialized service provision by professionals and at the same time by a quick expansion of knowledge and information, opportunities to exploit informational asymmetries at the expense of consumers who either do not have the resources or the time to obtain this information are likely to increase.

At the same time, improved access to information and the development of new digital technologies can help improve the position of consumers and restrain fraud. Passengers taking taxi rides in an unknown city are no longer constrained to remaining uninformed, as long as they have access to a mobile phone with a GPS function. Patients have better access to information about diseases and treatments than ever before, and nowadays there exist numerous platforms that offer medical advice online, either by doctors or by robots using sophisticated algorithms – noting, however, the caveats that are associated with relying on medical information or advice obtained from dubious sources on the internet. Similar kinds of tools are available to customers seeking advice regarding financial, legal or insurance services. Besides access to general information about products or services, consumers

have increasing opportunities to learn about the service provision of individual experts, for instance by taking advantage of online review platforms. This can be a powerful tool to reduce fraud and mistreatment through the beneficial effects of competition and reputation. We believe it is important to shed light into the way that digitalization and new technologies are likely to affect credence goods in the present and future and hold this to be an exciting and promising avenue for future research.

References

- Agarwal, R., Liu, C., Prasad, K., 2019. Personal search, second opinions, and the diagnostic effort of experts. *J. Econ. Behav. Organ.* 158, 44–61.
- Alger, I., Salanie, F., 2006. A theory of fraud and overtreatment in experts markets. *J. Econ. Manag.* 15 (4), 853–881.
- Anagol, S., Cole, S., Sarkar, S., 2017. Understanding the advice of commission-motivated agents: Evidence from the Indian life insurance market. *Rev. Econ. Stat.* 99 (1), 1–15.
- Arrow, K., 1963. Uncertainty and the welfare economics of medical care. *Amer. Econ. Rev.* 53, 941–973.
- Baker, L.C., 2010. Acquisition of MRI equipment by doctors drives up imaging use and spending. *Health Aff.* 29 (12), 2252–2259.
- Baksi, S., Bose, P., 2007. Credence goods, efficient labelling policies, and regulatory enforcement. *Environ. Resour. Econ.* 37 (2), 411–430.
- Balafoutas, L., Beck, A., Kerschbamer, R., Sutter, M., 2013. What drives taxi drivers? A field experiment on fraud in a market for credence goods. *Rev. Econom. Stud.* 80, 876–891.
- Balafoutas, L., Fornwagner, H., Kerschbamer, R., Sutter, M., Tverdostup, M., 2020. Diagnostic Uncertainty and Insurance in Credence Goods Markets. Mimeo, Department of Public Finance, University of Innsbruck.
- Balafoutas, L., Kerschbamer, R., Sutter, M., 2015. The hidden costs of tax evasion – collaborative tax evasion in markets for expert services. *J. Publ. Econ.* 129, 14–25.
- Balafoutas, L., Kerschbamer, R., Sutter, M., 2017. Second-degree moral hazard in a real-world credence goods market. *Econom. J.* 127, 1–18.
- Baron, D.P., 2011. Credence attributes, voluntary organizations, and social pressure. *J. Publ. Econ.* 95 (11), 1331–1338.
- Beck, A., Kerschbamer, R., Qiu, J., Sutter, M., 2013. Shaping beliefs in experimental markets for expert services: Guilt aversion and the impact of promises and money-burning options. *Games Econ. Behav.* 81, 145–164.
- Beck, A., Kerschbamer, R., Qiu, J., Sutter, M., 2014. Car mechanics in the lab – Investigating the behavior of real experts on experimental markets for credence goods. *J. Econ. Behav. Organ.* 108, 166–173.
- Bester, H., Dahm, J., 2018. Credence goods, costly diagnosis, and subjective evaluation. *Econom. J.* 128 (611), 1367–1394.
- Bonroy, O., Lemarie, S., Tropeano, J.P., 2013. Credence goods, experts and risk aversion. *Econom. Lett.* 120 (3), 464–467.
- Brenndorf, V., Kübler, D., Normann, H.T., 2015. Privacy concerns, voluntary disclosure of information, and unraveling: An experiment. *Eur. Econ. Rev.* 75, 43–59.
- Brosig-Koch, J., Hennig-Schmidt, K., Kairies-Schwarz, N., Wiesen, D., 2016. Using artefactual field and lab experiments to investigate how fee-for-service and capitation affect medical service provision. *J. Econ. Behav. Organ.* 131 (B), 17–23.
- Bucher-Koenen, T., Koenen, J., 2015. Do Seemingly Smarter Consumers Get Better Advice? Max Planck Institute for Social Law and Social Policy Discussion Paper No. 01-2015.
- Cason, T., Gangadharan, L., 2002. Environmental labeling and incomplete consumer information in laboratory markets. *J. Environ. Econ. Manag.* 43, 113–134.
- Chalkley, M., Malcomson, J., 1998. Contracting for health services when patient demand does not reflect quality. *J. Health Econ.* 17, 1–19.
- Clemens, J., Gottlieb, J.D., 2014. Do physicians' financial incentives affect treatment patterns and patient health? *Amer. Econ. Rev.* 104 (4), 1320–1349.
- Currie, J., Lin, W., Zhang, W., 2011. Patient knowledge and antibiotic abuse: evidence from an audit study in China. *J. Health Econ.* 3, 933–949.
- Darby, M.R., Karni, E., 1973. Free competition and the optimal amount of fraud. *J. Law Econ.* 16 (1), 67–88.
- Das, J., Holla, A., Mohpal, A., Muralidharan, K., 2016. Quality and accountability in healthcare delivery: Audit-study evidence from primary care in India. *Amer. Econ. Rev.* 106 (12), 3765–3799.
- De Jaegher, K., 2012. The value of private patient information in the physician-patient relationship: a game-theoretic account. *Comput. Math. Methods Med.* 1–16.
- Domenighetti, G., Casabianca, A., Gutzwiller, F., Martinoli, S., 1993. Revisiting the most informed consumer of surgical services: the physician-patient. *Int. J. Technol. Assess. Health Care* 9 (4), 505–513.

- Dranove, D., Jin, G.Z., 2010. Quality disclosure and certification: Theory and practice. *J. Econ. Lit.* 48 (4), 935–963.
- Dulleck, U., Gong, J., Li, J., 2013. Contracting for infrastructure projects as credence goods. *J. Publ. Econ. Theory* 17 (3), 328–345.
- Dulleck, U., Johnston, D., Kerschbamer, R., Sutter, M., 2012. The Good, the Bad and the Naive: Do Fair Prices Signal Good Types or Do They Induce Good Behaviour? IZA Discussion Paper No. 6491.
- Dulleck, U., Kerschbamer, R., 2006. On doctors, mechanics, and computer specialists: The economics of credence goods. *J. Econ. Lit.* 44 (1), 5–42.
- Dulleck, U., Kerschbamer, R., 2009. Experts vs. discounters: Consumer free-riding and experts withholding advice in markets for credence goods. *Int. J. Ind. Organ.* 27 (1), 15–23.
- Dulleck, U., Kerschbamer, R., Sutter, M., 2011. The economics of credence goods: An experiment on the role of liability, verifiability, reputation, and competition. *Amer. Econ. Rev.* 101, 530–559.
- Dulleck, U., Wigger, B.U., 2015. Politicians as experts, electoral control, and fiscal restraints. *J. Publ. Econ.* 121, 106–116.
- Ely, J., Fudenberg, D., Levine, D., 2008. When is reputation bad? *Games Econ. Behav.* 63, 498–526.
- Ely, J., Välimäki, J., 2003. Bad reputation. *Q. J. Econ.* 118 (3), 785–814.
- Etilé, F., Teyssier, S., 2016. Signaling corporate social responsibility: Third-party certification vs. brands. *Scand. J. Econ.* 118 (3), 397–432.
- Farhi, E., Lerner, J., Tirole, J., 2013. Fear of rejection? Tiered certification and transparency. *Rand J. Econ.* 44 (4), 610–631.
- Fedderson, T.J., Gilligan, T.W., 2001. Saints and markets: Activists and the supply of credence goods. *J. Econ. Manag. Strategy* 10 (1), 149–171.
- Fischer, C., Lyon, T.P., 2014. Competing environmental labels. *J. Econ. Manag. Strategy* 23 (3), 692–716.
- Fong, Y., 2005. When do experts cheat and whom do they target. *Rand J. Econ.* 36, 113–130.
- Fong, Y., Liu, T., Wright, D.J., 2014. On the role of verifiability and commitment in credence goods markets. *Int. J. Ind. Organ.* 37, 118–129.
- Frankel, A., Schwarz, M., 2014. Experts and their records. *Econ. Inq.* 52 (1), 56–71.
- Glaeser, E., Ujhelyi, G., 2010. Regulating misinformation. *J. Publ. Econ.* 94 (3–4), 247–257.
- Godager, G., Wiesen, D., 2013. Profit or patients' health benefit? Exploring the heterogeneity in physician altruism. *J. Health Econ.* 32, 1105–1116.
- Gottschalk, F., Mimra, W., Waibel, C., 2020. Health services as credence goods: A field experiment. *Econom. J.* (forthcoming).
- Green, E.P., 2014. Payment systems in the healthcare industry: An experimental study of physician incentives. *J. Econ. Behav. Organ.* 106, 367–378.
- Greiner, B., Zhang, L., Tang, C., 2017. Separation of prescription and treatment in health care markets: A laboratory experiment. *Health Econ.* 26 (S3), 21–35.
- Grosskopf, B., Sarin, R., 2010. Is reputation good or bad? An experiment. *Amer. Econ. Rev.* 100 (5), 2187–2204.
- Gruber, J., Owings, M., 1994. Physician financial incentives and cesarean section delivery. *Rand J. Econ.* 27 (1), 99–123.
- Harbaugh, R., Maxwell, J.W., Roussillon, B., 2011. Label confusion: The groucho effect of uncertain standards. *Manage. Sci.* 57 (9), 1512–1527.
- Hennig-Schmidt, H., Selten, R., Wiesen, D., 2011. How payment systems affect physicians' provision behavior – An experimental investigation. *J. Health Econ.* 30, 933–949.
- Hennig-Schmidt, H., Wiesen, D., 2014. Other-regarding behavior and motivation in health care provision: An experiment with medical and non-medical students. *Soc. Sci. Med.* 108, 156–165.
- Hilger, N., 2016. Why don't consumers trust experts? Credence goods with unobserved cost functions. *J. Law Econ.* 59 (2), 293–311.
- Huck, S., Luenser, G., Tyran, J.R., 2012. Competition fosters trust. *Games Econ. Behav.* 76 (1), 195–209.
- Huck, S., Luenser, G., Tyran, J.R., 2016a. Price competition and reputation in markets for experience goods: An experimental study. *Rand J. Econ.* 47, 99–117.
- Huck, S., Lünser, G., Spitzer, F., Tyran, J.R., 2016b. Medical insurance and free choice of physician shape patient overtreatment: A laboratory experiment. *J. Econ. Behav. Organ.* 131 (B), 78–105.
- Hyndman, K., Ozerturk, S., 2011. Consumer information in a market for expert services. *J. Econ. Behav. Organ.* 80 (3), 628–640.
- Iizuka, T., 2007. Experts' agency problems: Evidence from the prescription drug market in Japan. *Rand J. Econ.* 38 (3), 844–862.
- Inderst, R., Ottaviani, M., 2009. Misselling through agents. *Amer. Econ. Rev.* 99 (3), 883–908.
- Inderst, R., Ottaviani, M., 2012a. Competition through commissions and kickbacks. *Amer. Econ. Rev.* 102 (2), 780–809.
- Inderst, R., Ottaviani, M., 2012b. Financial advice. *J. Econ. Lit.* 50 (2), 494–512.
- Inderst, R., Ottaviani, M., 2012c. How (not) to pay for advice: A framework for consumer financial protection. *J. Financ. Econ.* 105 (2), 393–411.
- Jack, W., 2005. Purchasing health care services from providers with unknown altruism. *J. Health Econ.* 24, 73–93.
- Kerschbamer, R., 2018. Credence Goods in the Literature: Part 1 – Theory and Lab Experiments. Handout for the Innsbruck Winterschool on Credence Goods, Incentives and Behavior. Mimeo, Department of Economics, University of Innsbruck.
- Kerschbamer, R., Neururer, D., Sutter, M., 2016. Insurance coverage of customers induces dishonesty of sellers in markets for credence goods. *Proc. Natl. Acad. Sci.* 113 (27), 7454–7458.
- Kerschbamer, R., Neururer, D., Sutter, M., 2019. Credence Goods Markets and the Informational Value of New Media: A Natural Field Experiment, Working Papers in Economics and Statistics 2019-02. University of Innsbruck.
- Kerschbamer, R., Sutter, M., 2017. The economics of credence goods – A survey of recent lab and field experiments. *CESifo Econ. Stud.* 63, 1–23.
- Kerschbamer, R., Sutter, M., Dulleck, U., 2017. How social preferences shape incentives in (experimental) markets for credence goods. *Econom. J.* 127 (600), 393–416.
- Kreps, D.M., Wilson, R., 1982. Reputation and imperfect information. *J. Econom. Theory* 27 (2), 253–279.
- Liu, T., 2011. Credence goods markets with conscientious and selfish experts. *Internat. Econom. Rev.* 52, 227–244.
- Liu, T., Ma, C., 2013. Health insurance, treatment plan, and delegation to altruistic physician. *J. Econ. Behav. Organ.* 85, 79–96.
- Loewenstein, G., Cain, D.M., Sah, S., 2011. The limits of transparency: Pitfalls and potential of disclosing conflicts of interest. *Amer. Econ. Rev.* 101 (3), 423–428.
- Lu, F., 2014. Insurance coverage and agency problems in doctor prescriptions: Evidence from a field experiment in China. *J. Dev. Econ.* 106 (1), 156–167.
- Lyon, T.P., Maxwell, J.W., 2011. Greenwash: Corporate environmental disclosure under threat of audit. *J. Econ. Manag. Strategy* 20 (1), 3–41.
- Makris, M., Siciliani, L., 2013. Optimal incentive schemes for altruistic providers. *J. Publ. Econ. Theory* 15 (5), 675–699.
- Mimra, W., Rasch, A., Waibel, C., 2016a. Price competition and reputation in credence goods markets: Experimental evidence. *Games Econ. Behav.* 100, 337–352.
- Mimra, W., Rasch, A., Waibel, C., 2016b. Second opinions in markets for expert services: Experimental evidence. *J. Econ. Behav. Organ.* 131, 106–125.
- Mullainathan, S., Noeth, M., Schoar, A., 2012. The Market for Financial Advice: An Audit Study. NBER Working Paper No. 17929.
- Pesendorfer, W., Wolinsky, A., 2003. Second opinions and price competition: Inefficiency in the market for expert advice. *Rev. Econom. Stud.* 70 (2), 417–437.
- Rasch, A., Waibel, C., 2018. What drives fraud in a credence goods market?—Evidence from a field study. *Oxford Bull. Econ. Stat.* 80 (3), 605–624.
- Schneider, H., 2012. Agency problems and reputation in expert services: Evidence from auto repair. *J. Ind. Econ.* 60, 406–433.
- Schneider, T., Meub, L., Bizer, K., 2016. Consumer Information in a Market for Expert Services: Experimental Evidence. CEGE - Center for European, Governance and Economic Development Research Discussion Paper No. 285.
- Sülzle, K., Wambach, A., 2005. Insurance in a market for credence goods. *J. Risk Insurance* 72, 159–176.
- Wolinsky, A., 1995. Competition in markets for credence goods. *J. Inst. Theor. Econ.* 151 (1), 117–131.