



The hidden costs of tax evasion. Collaborative tax evasion in markets for expert services



Loukas Balafoutas^a, Adrian Beck^b, Rudolf Kerschbamer^{b,*}, Matthias Sutter^{c,a}

^a Department of Public Finance, University of Innsbruck, Austria

^b Department of Economics, University of Innsbruck, Austria

^c Department of Economics, University of Cologne, Germany

ARTICLE INFO

Article history:

Received 26 January 2014

Received in revised form 29 May 2015

Accepted 8 June 2015

Available online 17 June 2015

JEL classification:

C72

C91

D82

H26

Keywords:

Credence goods

Expert services

Tax evasion

Fraud

Experiment

ABSTRACT

We experimentally examine the impact of tax evasion attempts on the performance of credence goods markets, where contractual incompleteness results from asymmetric information on the welfare maximizing quality of the goods. Our results suggest that tax evasion attempts – independently of whether they are successful or not – lead to efficiency losses in the form of too low quality and less frequent trade. Thus, shadow economies may reduce welfare not only by inducing agents to incur costs to hide or to uncover taxable transactions, by imposing risk on uncertainty-averse tax evaders and by distorting competition, but also by creating an additional efficiency loss in the underlying market by forfeiting possible gains from trade and by inducing insufficient quality provision. We call these the hidden costs of tax evasion.

© 2015 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/>).

1. Introduction

Each year trillions of dollars slip through the tax authorities' fingers as a consequence of tax evasion. In the United States, for example, the overall net tax gap in 2006 was estimated to be approximately \$385 billion (IRS, 2012). A substantial body of theoretical and empirical work – briefly discussed at the end of the introduction – addresses the problem, focusing mainly on the determinants of the occurrence and the magnitude of tax evasion and on its welfare consequences. Considering the government as a provider of public goods and services financed by tax revenues, tax evasion is detrimental to welfare for the simple reason that it adds to the excess burden of taxation because some of the costs of evasion are real resource costs and not just transfers (Feldstein, 1995, 1999; Chetty, 2009).¹ For instance, a real resource cost can

emerge when taxpayers try to conceal and tax authorities try to detect tax evasion (see, e.g., Bayer, 2006), when tax evasion imposes uncertainty on risk-averse evaders (Yitzhaki, 1987) and when tax-evading firms drive tax-honest ones out of the market (Strand, 2005). In this paper we argue that in markets where transactions are governed by contractual incompleteness an additional welfare impact of attempted and successful tax evasion might result from its effect on agents' moral incentives of exploiting their informational advantage to their benefit and to the detriment of their trading partners, hence leading to a substantial drop in overall efficiency.

¹ The standard thought experiment for the measurement of the excess burden of taxation is to calculate the net loss from raising the tax rate and returning the revenue lump sum to the taxpayer. As Feldstein (1995, 1999) first asserted, under some conditions the 'tax-base-elasticity' (for the case of income taxation better known as the 'elasticity of taxable income') is a sufficient statistic for the marginal welfare loss of raising the tax rate (see Saez et al., 2012, for a discussion). The tax-base-elasticity approach implicitly assumes that – at the agent's optimal solution – the marginal social cost of reducing the tax base by a dollar equals the tax rate for all behavioral responses intended to reduce the burden of taxation. As Slemrod (1998) and Chetty (2009) have convincingly argued, for tax evasion as the behavioral response this assumption is likely to be violated in practice because some of the private costs of evasion are transfers rather than real resource costs.

* Corresponding author at: Department of Economics, University of Innsbruck, Universitätsstrasse 15, A-6020 Innsbruck, Austria. Tel.: +43 512 507 7400.
E-mail address: Rudolf.Kerschbamer@uibk.ac.at (R. Kerschbamer).

Key to our argument is the observation that in many markets the approval of both trading partners is needed to evade taxes.² This cooperative dimension of tax evasion implies that at some point the trading partners have to reveal their preferences regarding tax evasion to each other. In markets governed by incomplete contracts it seems plausible that the revelation of an agent's attitude regarding tax evasion influences the trading partner's sentiments towards the agent and therewith the behavior in the underlying market.

We will experimentally investigate the impact of revealed intentions regarding tax evasion on the performance of a market characterized by incomplete contracts — a topic that has to the best of our knowledge not yet been explored in the literature. The setting in our controlled laboratory experiment is a market for expert services — often also referred to as a *credence goods* market (see Dulleck and Kerschbamer, 2006, for a general framework and an overview of the theoretical literature). Such markets are plagued by serious contractual incompleteness due to superior information of the seller on the quality that yields the highest surplus from trade. Prime examples are health services, where the doctor knows better than the patient which disease the latter has and which treatment is appropriate, and all kinds of repair services, where the expert knows more about the type of service the item needs than the owner. On top of the asymmetric information on the welfare maximizing quality, in many markets for expert services there is also asymmetric information on the quality provided. For instance, in the market for medical services a patient might be unable to tell ex post whether the injection he received contained a high-cost substance or not. The informational asymmetries imply that complete contracts are infeasible and open the door to a whole array of different types of fraudulent behavior on the sellers' side, including overtreatment (providing a higher quality than the surplus maximizing one), undertreatment (choosing a quality that is insufficient to satisfy the consumer's needs), and overcharging (charging for a higher quality than has been provided). Furthermore, anticipated fraud might lead consumers to abstain from the market altogether, leading to further efficiency losses.³

Our main hypothesis is that the size of the efficiency losses resulting from contractual incompleteness depends on the trading partners' revealed intentions to evade taxes. On the one hand, a mutual agreement to pay or to evade taxes can decrease the social distance between the trading parties. Reduced social distance has been demonstrated to lead to more cooperative behavior, for instance by Charness et al. (2007) and Götte et al. (2012). Based on this evidence we would expect that agents behave nicer towards trading partners who reveal the same attitude regarding tax evasion. In the case where a mutual agreement to *pay* taxes leads to nicer (i.e., more trustful or more trustworthy) behavior we refer to this as the *solidarity effect*, while in the case where a mutual agreement to *evade* taxes leads to nicer behavior we refer to a *conspiracy effect*.

On the other hand, in constellations where one trading partner reveals the intention to pay while the other reveals the intention to evade taxes each player might have negative emotions towards the other, leading to less trustful and trustworthy behavior and thereby decreasing the efficiency of a transaction. We call this the *punishment effect*. Systematic differences in the market behavior of sellers in different constellations of revealed intentions to evade taxes might

also arise if the revealed intentions to evade taxes are driven by preferences that also drive market behavior. For instance, it seems plausible that more pro-social sellers are not only less inclined to evade taxes but also less willing to defraud consumers.⁴ We call this the *selection effect*. Finally, differences in behavior may also be due to the sheer fact that taxes are actually evaded in one but paid in another constellation, which directly affects the resulting profits from the transaction. We call this the *cake size effect*. In addition to these direct effects an impact on efficiency might also arise if sellers who reveal their willingness to pay or to evade taxes are expected to be more or less likely to defraud their consumers (because of the solidarity, conspiracy, punishment, selection or cake size effect). Consumers may then be more or less willing to interact, producing a further impact on efficiency. Given all those effects it is not a priori clear whether attempted or actual tax evasion has a positive, negative, or no impact on the efficiency of the underlying market.

To assess the empirical relevance of these effects we ran an experiment with 248 subjects who were either in the role of a seller or a consumer on a credence goods market. Our experiment involved three treatments implemented in a between-subjects design. In our main treatment, before any interaction could take place, both sellers and consumers had to indicate whether they wanted to evade or to pay taxes. Only when both parties to a potential trade agreed, taxes were evaded, otherwise they were paid.⁵

The results of our main treatment suggest that revealed intentions to evade taxes lead to efficiency losses of up to 50% of the available surplus, independently of whether taxes are actually evaded or not. In the experiment the drop in efficiency is driven by large differences in interaction and undertreatment rates of up to 20 percentage points. In particular, tax-honest consumers are much less likely to enter the credence goods market if they know they will have to interact with a tax-dishonest seller. Also, if a transaction takes place, the undertreatment rate is higher when at least one of the agents has revealed the intention to evade taxes than in constellations where both transaction partners are tax honest.

Using control treatments to disentangle the various effects of revealed intentions regarding tax compliance, we find evidence for the existence of a solidarity effect between two tax-honest agents and a conspiracy effect between two tax-dishonest agents, both leading to less fraud when agents reveal the same tax attitude. We also find some limited evidence for the existence of a selection effect, with tax-dishonest sellers being more likely to defraud their customers, but almost no evidence for the existence of a punishment effect or a cake size effect.

Overall, our results suggest that tax evasion may reduce welfare not only by inducing agents to incur costs to hide or to uncover taxable transactions, by imposing risk on risk-averse tax evaders and by distorting competition between tax-honest and tax-evading firms, but also by creating an additional ('hidden') efficiency loss in the underlying

² This is not the case for all bilateral transactions, of course. For instance, in the case of sales taxes or excise taxes, the seller may provide a false invoice to the buyer who remains unaware that the tax is evaded. While we will also study a situation where tax evasion can be imposed unilaterally by the seller, our focus will be on transactions where both partners must agree to evade taxes.

³ Although standard theory assuming own-money-maximizing and risk-neutral preferences predicts low efficiency in markets governed by incomplete contracts, recent experimental studies — by Dulleck et al. (2011) and Beck et al. (2013) for markets for credence goods and by Huck et al. (2010, 2012, 2013) for markets for experience goods — have shown that such markets work considerably better than predicted, probably due to agents' social preferences, as examined in Kerschbamer et al. (forthcoming).

⁴ While in the context of our experiment — where the tax receipts are contributed to a public good which benefits all subjects — the link between being a tax evader and being an anti-social agent seems natural, this link might be less apparent in reality. Indeed, as one of the anonymous referees pointed out, in countries where the government is highly corrupted, over-regulates the market, or imposes too strong fiscal pressure on citizens, tax evasion might be a natural, ethically acceptable behavior — see Schneider and Enste (2000, p.108) for a discussion. This raises the question of how important the phenomena studied here are for different markets. Experimental evidence presented by Güth et al. (2005) suggests that tax evasion decisions do not depend on whether tax receipts are redistributed with overall efficiency gains or losses (the latter mimicking waste or corruption in government). Whether this 'neutrality' result extends to the present framework is an open question and we consider it as important to examine this question in more depth in future studies.

⁵ Note that in real markets differences in behavior between the taxed and the untaxed sector might also arise from differences in institutional consumer protection. For instance, liability rules might only apply in the taxed but not in the untaxed sector. We abstract from such differences to isolate the pure effect of collaborative tax evasion on the performance of the market in which the transaction takes place.

market through less frequent trade and more frequent provision of insufficient quality in case of trade. While our work provides a first indication of the existence of such a hidden cost of tax evasion, more work is needed to obtain information on its determinants and empirical importance. Our hope is that the present paper stimulates further studies in the lab and the field addressing those questions.

Before turning to our game and the experiment we refer to related literature, of which several strands are relevant. First, our study is related to the literature on the economics of tax evasion in general (Andreoni et al., 1998; Slemrod, 2007). A substantial body of theoretical and empirical work addresses the problem, focusing mainly on the effects of institutional remedies against tax evasion. The seminal work by Allingham and Sandmo (1972) models citizens' decisions about which fraction of income to declare as a simple portfolio decision problem and determines the effects of changes in exogenous factors such as tax and penalty rates and the probability of detection. Later work – by Yitzhaki (1979), Reinganum and Wilde (1985), and Alm et al. (1992), for instance – refines the analysis of optimal taxation, again mainly focusing on the effects of changes in the institutional framework in which tax evasion potentially occurs. The theoretical work is accompanied by empirical studies analyzing field data (see, e.g., Clotfelter, 1983; Feinstein, 1991; Slemrod, 2007) as well as by laboratory and field experiments (Baldry, 1986, 1987; Torgler, 2002, 2004; Kleven et al., 2011). None of these studies has investigated how attitudes towards tax evasion might affect the behavior in markets with incomplete contracts.

More closely related to the present paper is the (scarce) literature on the cooperative dimension of tax evasion. Early contributions are the theoretical studies by Yaniv (1992), Boadway et al. (2002), and Chang and Lai (2004), mainly focusing on the effects of exogenous changes in either the probability of detection or the penalty in case of detection on collaborative tax evasion. More recently, Kleven et al. (2009) show in an agency model that collaborative tax evasion is fragile in the presence of verifiable business records and many involved parties; and Madzharova (2013) and Naritomi (2013) investigate empirically the effects of changing the incentives for collusive tax evasion on the producer side (by changing the difference between the payroll tax rate and the corporate income tax rate), or the consumer side (by giving monetary rewards to consumers to ask for receipts). None of these studies addresses the question how revealed intentions to pay or evade taxes might affect agents' moral incentives of exploiting their informational advantage in markets governed by contractual incompleteness.

Regarding moral incentives in tax-related decisions our study is related to the literature studying the impact of non-standard preferences on tax evasion decisions. For instance, Spicer and Becker (1980) study the impact of fairness perceptions on this decision, Gordon (1989) introduces a psychological cost of evasion, Erard and Feinstein (1994) adapt the concepts of guilt and shame to the context of tax compliance, Frey and Torgler (2007) study the impact of perceptions regarding the compliance of other taxpayers on evasion, and Konrad and Qari (2012) study the effect of patriotism on compliance. By contrast, our focus is not on the compliance decision per se but rather on the question whether different outcomes on the compliance stage have different implications for the subsequent behavior of agents trading in the market and thus for the efficiency of that market.

2. A credence goods game

2.1. Basic framework

We base our experimental investigation on a simplified version of a model on expert services by Dulleck and Kerschbamer (2006). In this two-player sequential-move game, a consumer (c ; she) has a problem θ which is either major and requires a high-quality service, or minor and requires a low-quality service. While the consumer herself only knows the ex-ante probability h ($1 - h$) of having the major (minor) problem, an expert seller (s ; he) receives full information about the

type of the consumer's problem. More specifically, the sequence of the game is as follows: After observing exogenously given prices p^l (representing the price of a low-quality service) and p^h (representing the price of a high-quality service), with $p^l < p^h$, the consumer decides whether to buy a service from the expert seller or to opt out of the market. In the former case, a random move by nature (n) determines the consumer's problem; the seller learns the realization of this random move (by performing a free diagnosis) and then provides a service of either low (l) or high (h) quality. We denote the index of the quality provided by $\tau \in \{l, h\}$. The high quality service ($\tau = h$) solves both types of problems, while the low quality service ($\tau = l$) is only sufficient for minor problems. Let c^τ denote the costs of quality τ , and $c^l < c^h$. For the quality he claims to have provided the expert can then charge the associated price (either p^l or p^h). We denote the index of the quality charged for by $\iota \in \{l, h\}$.

The material payoffs in case of an interaction are such that the seller obtains the price he charged from the consumer and has to bear the costs for the quality provided, i.e., in the baseline game the seller's profit is $\pi_s^b = p^\iota - c^\tau$. For future reference we define a situation in which the expert provides the low-quality service and charges for the high-quality service as *overcharging*; and we refer to a situation in which the expert provides the low quality to a consumer who has the major problem as *undertreatment*.⁶ The consumer's payoff in case of an interaction depends on whether the provided service solved her problem: She receives the value $V > 0$ if the quality was sufficient ($\tau \in \{l, h\}$ if $\theta = l$; and $\tau = h$ if $\theta = h$) and zero otherwise; and she has to pay the price charged by the seller in any case. That is, $\pi_c^b = V - p^\iota$ if τ is sufficient, and $\pi_c^b = -p^\iota$ if it is not. In case of no interaction both the consumer and the expert receive an outside value of o . Fig. 1 presents the structure of this baseline game, excluding the price announcement stage.

2.2. Tax evasion

Taking the baseline game in Fig. 1 as starting point, we now modify it in two ways to implement a tax evasion framework. More specifically, we allow the baseline game to be played either with or without taxes. In the *taxation variant of the game* (denoted Γ_t), the monetary payoffs of the two agents change to $\pi_i^t = \pi_i^b - t$ for $i = s, c$. For the sake of simplicity, the tax is modeled as a flat payment t imposed on each side of the market whenever a transaction takes place no matter which quality is provided and which price is charged for the good. In the *evasion variant of the game* (denoted Γ_e) the material payoffs of the two players depend upon whether tax evasion is detected or not. In the latter case $\pi_i^e = \pi_i^b$, while in the former case $\pi_i^e = \pi_i^b - (t + f)$, where f is the fine levied by the tax-authority and $i = s, c$. Whether tax evasion is detected or not depends on the realization of an independent random variable, with the probability of detection being d . The sum of all tax payments and penalties levied are multiplied with a factor $r > 1$ and are then redistributed equally among all market participants in the manner of a public goods game. To preserve the public good nature of the structure of material payoffs the factor r is chosen such that the individual return of contributing is less than one (i.e., $r/n < 1$, where n is the number of market participants).

In order to detect a possible relationship between the revealed intention of an agent to pay or to evade taxes and the way she behaves

⁶ In theory there exists a third type of fraudulent behavior, which is *overtreatment*. It occurs whenever the expert provides the high quality to a consumer who has the minor problem. When quality is not observed by the consumer – as it is the case in the games we study in this paper – overtreatment is dominated for the seller in material terms by overcharging. Taking into account moral incentives it might still become an issue if it is morally less costly for the seller to charge the high price when the high quality service is provided (even if the consumer does not need it). However, as the results presented by Dulleck et al. (2011) suggest, this theoretical possibility is irrelevant in an experiment. This is confirmed in our results: From the 2480 cases where overtreatment is possible (because the customer has the minor problem) the high quality is provided in only 58 cases.

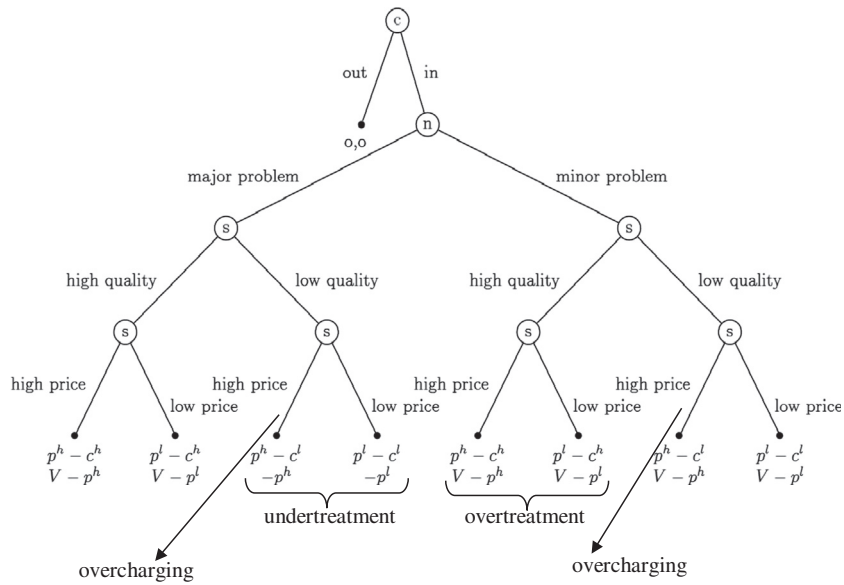


Fig. 1. Structure of the baseline game representing a credence good situation. Note. The term *undertreatment* refers to providing the low quality to a consumer who needs the high quality; *overtreatment* refers to providing the high quality to a consumer who needs the low quality; and *overcharging* refers to charging the high price when the low quality has been provided.

or the way she is treated by another agent in the market stage of the game we introduce a tax-decision stage in which the seller and his consumer can jointly determine whether they want to play variant Γ_e or variant Γ_t of the market game (see Fig. 2). We accomplish this by giving the expert and the consumer the possibility of simultaneously and independently expressing their willingness to evade taxes as a binary decision. In what follows the term “tax-dishonest” is used to indicate that a player revealed the intention to evade

taxes at this stage, while “tax-honest” means that he or she decided against tax evasion. It is reasonable to assume that no one can be forced to evade taxes. Therefore, we require that tax evasion needs a unanimous consent, i.e., taxes are evaded if and only if both parties of a transaction are tax-dishonest in the tax-decision stage of the game. After this stage, both agents get to know the tax decision of the other player and continue with playing either the Γ_e or the Γ_t variant of the market game.

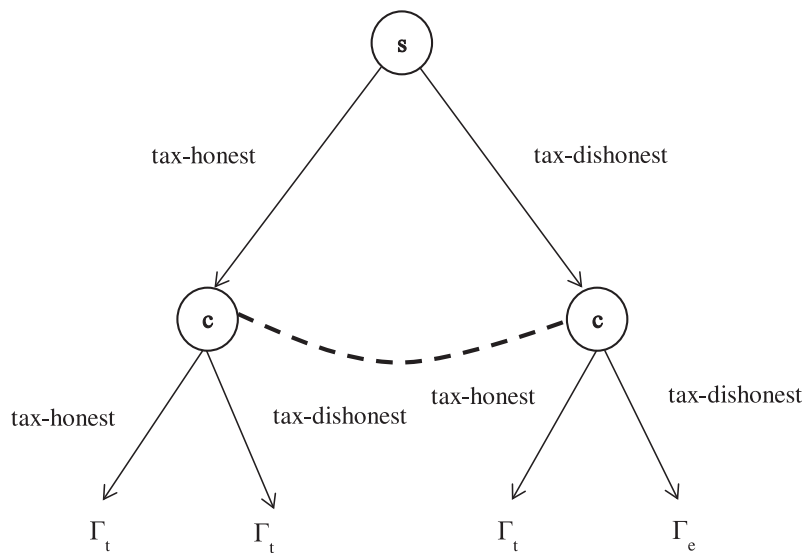


Fig. 2. Tax-decision stage. Note. The seller (s) and the buyer (c) simultaneously state whether they want to pay taxes (“tax-honest”) or evade taxes (“tax-dishonest”). If at least one player wants to pay taxes, game Γ_t is played subsequently. If both players want to evade taxes, game Γ_e is played.

Assuming common knowledge that all agents are risk-neutral and exclusively interested in their own material payoff and implementing parameters such that $(1 - h)V - p^h < o$, the subgame perfect prediction for the market game is that in case of a transaction the expert always provides the low-quality service and charges the price for the high-quality service. Anticipating this, the consumer opts out of the market. Expecting that the market will remain inactive, both parties of a transaction are indifferent between the two options at the tax-decision stage of the game.

3. Experimental design

3.1. Experimental treatments

In order to isolate the consequences of different constellations of revealed intentions regarding tax evasion on the behavior of the agents and to identify the causes for these consequences our experiment involves three treatments implemented in a between-subjects design.

In treatment *ENDO* (for *endogenous*) participants are asked to express their willingness to evade taxes as a binary decision, as described in the previous section. This is our main treatment, the findings of which we discuss in Section 4. In addition we run two control treatments, called *EXO* and *SELLER*.

In *EXO* (for *exogenous*) participants are randomly assigned to either the Γ_e or the Γ_t variant of the game. To control for unobserved factors, the parameters regarding price vectors, problem types, market conditions (taxed or untaxed), and detected instances of tax evasion are held constant across *EXO* and *ENDO*. This is done by taking the data observed in the *ENDO* treatment and implementing them exogenously as experimental parameters in the *EXO* treatment. We will use treatment *EXO* to examine whether differences in the market behavior of trading partners in different constellations of revealed preferences to evade taxes are caused by the shift in the market environment itself. It might be the case, for instance, that subjects behave differently in an environment where gains from trade are high with high probability and low with low probability (as it is the case in the Γ_e variant of the market game) as compared to an environment where gains are intermediate with certainty (as it is the case in the Γ_t variant of the game). In the Introduction we referred to this as the *cake size effect*.

In the *SELLER* treatment the power of deciding about tax evasion is entirely on the sellers' side, while consumers are not able to express their intentions regarding tax evasion. This treatment will be used to gain information about the existence of a *selection effect* (as discussed in the Introduction), in particular to shed light on the question whether tax-dishonest sellers are per se more likely to misbehave in credence goods markets. Table 1 summarizes our treatments.

3.2. Experimental procedures

At the beginning of each session subjects were randomly assigned to their fixed role as either a seller or a consumer and they were given an initial endowment of 200 experimental currency units (ecu), with 100 ecu being worth 1 euro. In all treatments subjects played 20 periods in a random stranger design within a matching group of eight participants, consisting of four sellers and four consumers. The random

stranger design was implemented in order to avoid reputation effects in the buyer–seller interaction. Although it is true that trading partners meet more than once on average so that some kind of reputation can emerge at the group level, subjects can never identify who their current partner is, thus excluding reputation formation at the individual level.

A consumer's probability of having the major problem was set to $h = 0.5$. We decided for equal probabilities of the two problems as this is particularly simple to explain to subjects and avoids confusion. The costs for the expert of providing a low-quality service was $c^l = 20$ ecu, the cost of a high-quality service was $c^h = 40$ ecu, and the value for the consumer of receiving a sufficient quality was $V = 120$ ecu. This parameter constellation implies relatively large efficiency gains for appropriate treatment and relatively low incentives for undertreatment (in comparison to the parameter constellation in Dulleck et al., 2011) and was chosen in the hope of getting relatively high participation rates and enough room for treatment variation.

Five price vectors (p^l, p^h) from the set $\{(50, 60), (50, 70), (50, 80), (60, 80), (70, 80)\}$ were imposed randomly and with equal frequency. We decided for variation in prices for two reasons. The first reason is that we wanted to avoid that subjects get annoyed, tired or bored when facing the same decision situation again and again. The second reason is to have a variation in material incentives for overcharging and – for those sellers who charge honestly – for under- and overtreatment.

The chosen set of price vectors satisfies certain desirable properties. First, it includes all types of price vectors – undertreatment price vectors [(50, 60) and (70, 80)], equal markup price vectors [(50, 70) and (60, 80)] and also an overtreatment price vector [(50, 80)]. Undertreatment (overtreatment) price vectors have the property that an honestly-charging expert has a material incentive to always provide the low (high) quality because $p^l - c^l > p^h - c^h$ ($p^l - c^l < p^h - c^h$, respectively). With equal markup vectors an honestly-charging expert has no incentive to provide the wrong quality, because $p^l - c^l = p^h - c^h$. Second, for the three price vectors (50, 60), (50, 70) and (50, 80), we fix the price for the low quality and vary the price for the high quality, while for the three price vectors (50, 80), (60, 80) and (70, 80) we fix the price of the high quality and vary the price for the low quality. This allows us to control for possible effects of changing one price while holding the other constant. Each of these five price vectors was played for four rounds, twice with the consumer having a minor problem and twice with the consumer having a major problem. The sequence of (price vector, problem)-pairs was randomized on the individual expert's level. The outside payment o was set to 5 ecu. With this value for o , the condition $(1 - h)V - p^h < o$ is fulfilled for all price constellations and standard theory predicts a complete market breakdown for our experimental market game.

Roughly in line with the audit probabilities for non-business individual returns and corporate returns in the US (cf. IRS, 2011) we set the probability of an untaxed transaction being detected to 2% ($d = 0.02$). To keep things simple in the experiment, the tax t as well as the penalty f was set to 10 ecu. All taxes and penalties were contributed to a public account with an efficiency factor of 2, i.e., all payments were doubled and equally distributed among the members of a matching group, implying a marginal per capita return of 0.25. Payments from the public good were only revealed at the end of the experiment.

At the outset of each session, participants were given plenty of time to read through the instructions (those for treatment *ENDO* are given in the Appendix A). Afterwards, the instructions were also read out aloud. The experiment was framed as a decision about tax evasion (and about problems and solutions in a market with buyers and sellers) in order to induce the underlying tax context rather than letting participants choose between abstract lotteries without context. Before the start of the experiment, each participant had to answer a set of control questions and the experiment did not continue until all participants had answered all questions correctly. To get the same number of data points

Table 1
Experimental treatments.

Treatment	Tax evaded if	N
<i>ENDO</i>	Both are tax-dishonest	80
<i>SELLER</i>	Seller is dishonest	88
<i>EXO</i>	Exogenously imposed	80

from all experts, an expert was asked to make his provision and charging decision after he received information about the tax attitude of his consumer but before he received information about the consumer's interaction decision.⁷

We concluded each session with a test eliciting each subject's risk attitudes – which could be an important factor driving tax evading behavior. Each subject was exposed to a series of 10 binary choices between a cash gamble and a safe payoff. While the cash gamble remained the same in all choices – it always yielded either 500 ecu or 0 ecu, each with 50% probability – the safe payoff increased in steps of 50 ecu from 50 ecu to 500 ecu. The 10 binary choices were presented to the subjects one-at-a-time in random order (i.e., each binary decision on an own screen). With the choice data we constructed a variable (called risk factor) defined as the number of times a subject opted for the risky lottery, with higher values thus corresponding to a higher willingness to take risks.⁸

The experiment was conducted computerized at the University of Innsbruck using z-Tree (Fischbacher, 2007) and recruitment was done via ORSEE (Greiner, 2004). In total 248 students participated, 80 in each of the treatments *ENDO* and *EXO* and 88 in treatment *SELLER*. Sessions lasted for about 75 min and average earnings were about 11 euros. Each subject took part in only one session and we excluded subjects with prior experience in credence goods games and tax experiments.

4. Results in *ENDO*

In what follows we define “e-sellers” as sellers who revealed the intention to evade taxes and “t-sellers” as sellers who revealed the intention to pay the tax, and apply the analogous definitions for consumers. We will refer to the pool of sellers and buyers who revealed a preference for paying (evading) the tax as “t-agents” (“e-agents”). For future reference we also define the following variables:

- The *tax-dishonesty rate* is the relative frequency of cases where subjects revealed the intention to evade taxes.
- The *interaction rate* (IR) is the relative frequency of consumers choosing to enter the market.
- The *undertreatment rate* (UR) is the ratio of cases where the seller chose to provide the low-quality service to a consumer having the major problem to all cases where the consumer has the major problem.
- The *overcharging rate* (OR) is the ratio of cases where the seller charged for the high-quality service while providing the low-quality service to all cases where a low-quality service was provided.⁹
- *Market efficiency* is defined as $IR(1 - UR)$, which allows comparing market performance after different outcomes of the tax-decision

⁷ Since the strategy method is not used to elicit a seller's provision and charging behavior contingent on the tax decision of the trading partner (which is the main focus of this work) but only for the participation decision, experimenter demand effects or similar objections related to the validity of the strategy method are arguably less of an issue in our design. For discussions of the pros and cons of using the strategy method see, for instance, Casari and Cason (2009), or Brandts and Charness (2011).

⁸ This measure of risk attitudes (which is similar to the one used by Holt and Laury, 2002) is problematic for subjects who make inconsistent choices (by switching more than once, switching in the wrong direction or choosing the gamble when it is first-order stochastically dominated). In our experiment only a small minority of subjects (6.45% of our sample) made inconsistent choices in the risk elicitation task. We have checked and confirmed the robustness of our main results by running versions of our regressions in which we excluded the inconsistent subjects.

⁹ It is important to note that the categories *undertreatment* and *overcharging* are not mutually exclusive. The category *overcharging*, in fact, comprises different forms of fraudulent behavior: fraud where there is “overcharging only” (that is, cases where the expert appropriately provides the low quality to a consumer who has the minor problem, but charges the high price for the low quality); and fraud where there is “overcharging and undertreatment” (that is, cases where the expert inappropriately provides the low quality to a consumer who has the major problem and in addition charges the high price for the low quality).

Table 2
Market behavior in treatments *ENDO*, *EXO*, and *SELLER*.

Situation	Seller	Consumer	Cond.	Freq.	UR	OR	IR	EFF
<i>Panel (a): ENDO</i>								
A (<i>ENDO</i>)	t-Seller	t-Consumer	Γ_t	0.17	0.43	0.93	0.85	0.48
B (<i>ENDO</i>)	t-Seller	e-Consumer	Γ_t	0.20	0.58	0.91	0.67	0.28
C (<i>ENDO</i>)	e-Seller	t-Consumer	Γ_t	0.26	0.65	0.91	0.64	0.23
D (<i>ENDO</i>)	e-Seller	e-Consumer	Γ_e	0.37	0.57	0.98	0.74	0.32
<i>Panel (b): SELLER</i>								
E (<i>SELLER</i>)	t-Seller	n/a	Γ_t	0.36	0.56	0.88	0.68	0.30
F (<i>SELLER</i>)	e-SELLER	n/a	Γ_e	0.64	0.70	0.97	0.59	0.18
<i>Panel (c): EXO</i>								
G (<i>EXO</i>)	Tax paid exogenously		Γ_t	0.63	0.59	0.92	0.71	0.29
H (<i>EXO</i>)	Tax evaded exogenously		Γ_e	0.37	0.62	0.92	0.72	0.27

Notes. Market behavior after different outcomes of the tax-decision stage. UR (OR) represents the undertreatment (overcharging) rate and IR the interaction rate. ‘freq.’ represents the frequency of occurrence of a certain situation, ‘cond.’ the resulting market condition (taxed or untaxed), and ‘EFF’ the resulting market efficiency calculated as $IR(1-UR)$.

stage of the game. Note that this expression represents the likelihood that the customer's problem is fixed and the value of $V = 120$ is created in a given matching – for this to happen the customer must interact and he must not be undertreated.

The focus of our interest lies in the consequences of different constellations of revealed intentions to pay or evade taxes in treatment *ENDO* on the subsequent behavior in the credence goods market. We therefore first report the results of this treatment and postpone the discussion of the two control treatments to the next section. As a reminder, the agreement of both trading partners was required to conduct tax evasion in treatment *ENDO*. Accordingly, we will refer to *attempted tax evasion* when one of the agents revealed the intention to evade taxes but the other did not, and we will refer to *actual tax evasion* when both agents revealed the intention to evade taxes so that the tax was not paid.

Table 2 reports the undertreatment and overcharging rates (i.e., provision and charging decisions by expert sellers) as well as interaction rates (i.e., entry decisions by consumers) and overall market efficiency, broken down by treatment and by configuration of agent types.¹⁰ As is clear from the table, the benchmark prediction of an inactive market (corresponding to an interaction rate of 0%) – and of an undertreatment and an overcharging rate of 100% in case of interaction – fails largely. The interaction rate is well above 50% in all the situations listed in the table and the undertreatment rate is between 43% (in situation A) and 70% (in situation F). The only variable that is roughly in line with the benchmark prediction is the overcharging rate – it varies between 88% in Situation E and 98% in Situation D and is therewith not far from the predicted 100%.

Comparing undertreatment and overcharging rates in the different situations reveals that the latter are much larger than the former in each row. This is in line with the findings in related experiments (see Dulleck et al., 2011, or Beck et al., 2013, for instance).¹¹ There are several possible explanations for this finding. An important one is based on the observation that undertreatment necessarily involves an efficiency loss

¹⁰ Table 2 does not report the results for overtreatment, given that overtreatment is costly from the seller's point of view and does not create any value for the consumer. As it turns out, overtreatment rates were indeed extremely low in all treatments and occurred overall in just 2.4% of cases with a minor problem.

¹¹ Among the 16 treatments investigated in Dulleck et al. (2011), the B/N constellation is closest to our credence goods game in terms of institutional and informational framework. For this treatment the authors report an undertreatment rate of 53% which is roughly in line with our results in Table 2. The overcharging rate is with 88% somewhat lower than in our experiment. A possible reason is that we provided an example in the instructions that features overcharging, while Dulleck et al. had no such example. In hindsight providing such an example was not a good idea because the observed high overcharging rates make it difficult to find statistically significant differences across situations.

while overcharging does not. Specifically, overcharging involves a pure (1:1) transfer of money from the customer to the expert. Thus, for an expert not to have an incentive to overcharge, the material payoff of the customer must receive at least the same weight in the expert's utility function as the expert's own material payoff – which is quite a demanding requirement. By contrast, undertreatment involves a transfer of money from the customer to the expert that is associated with an efficiency loss: If an expert seller undertreats his customer then his material payoff increases by $c^h - c^l$, while the payoff of the customer decreases by $V > c^h - c^l$. An implication is that a weight of $[(c^h - c^l)/V] < 1$ on the payoff of the customer is sufficient to prevent undertreatment, which is less demanding. A second explanation is based on the observation that undertreatment is observable with certainty by the customer, while overcharging is not (except for the case where it comes together with undertreatment). This is an important difference from the point of view of the expert, because the moral cost of fraudulent behavior is likely to be higher if the expert expects that the customer will realize ex post that she has been cheated, compared to a situation where the expert knows that the customer can never be sure whether a high price is due to overcharging or to nature having chosen the major problem.¹²

In discussing the effects of attempted and actual tax evasion on market outcomes we will proceed in reverse chronological order, following the logic of backwards induction: We discuss the provision and charging behavior of sellers in Subsection 4.1, the participation decision of consumers in 4.2, the overall market efficiency in 4.3 and the tax evasion decisions in 4.4.

4.1. The market game: undertreatment and overcharging

We begin by analyzing the effects of attempted and actual tax evasion on market outcomes and take a situation with two t-agents (Situation A in Table 2) as a reference. Situations B and C refer to attempted tax evasion by one of the two agents, while Situation D leads to actual tax evasion. Panel (a) of Table 2 reveals that the provision behavior of expert sellers is most honest when both agents have revealed their willingness to pay taxes in the tax-decision stage (Situation A). Here, the undertreatment rate of 0.43 is lower than in all other situations (and this value is significantly different from the corresponding values in B and in D, with $p < 0.05$ for each comparison).¹³ When being matched with an e-consumer, a t-seller on average significantly increases the UR from 0.43 to 0.58, while an e-seller decreases the UR from 0.65 to 0.57. However, this latter drop in UR is not significant. Also, comparing undertreatment in the cases where a seller and a consumer have agreed on evading the tax (Situation D) with the other situations reveals that undertreatment rates in Situation D do not differ significantly from those in situations B and C ($p > 0.8$), but are significantly higher than those in dyads in which both subjects were tax-honest ($p = 0.01$).

Columns (1) and (2) of Table 3 formally test the consequences of interacting with an e-consumer for the two types of sellers.¹⁴ Column (1) indicates that tax-dishonest consumers have to be prepared for a (weakly) significantly higher probability of undertreatment than honest

Table 3
Undertreatment rates and interaction rates, treatment ENDO.

Dependent variable	(1)	(2)	(3)	(4)
	Undertreatment	Undertreatment	Interaction	Interaction
Used data	t-Sellers	e-Sellers	t-Consumers	e-Consumers
Meeting an e-seller			-0.220*** (0.032)	0.080 (0.064)
Meeting an e-consumer	0.191* (0.107)	-0.099 (0.097)		
p^l	0.008 (0.006)	0.003 (0.004)	0.007* (0.004)	-0.007** (0.004)
p^h	-0.007 (0.007)	-0.008* (0.005)	-0.016** (0.008)	-0.009* (0.005)
Period	0.032*** (0.007)	0.011** (0.004)	-0.017*** (0.003)	-0.008* (0.004)
Female	-0.075 (0.020)	0.047 (0.130)	-0.118** (0.063)	0.003 (0.055)
Age	-0.031 (0.019)	-0.029 (0.018)	-0.006 (0.013)	-0.001 (0.008)
Risk factor	-0.010 (0.068)	-0.055 (0.076)	0.012 (0.036)	0.028 (0.025)
Observations	154	246	343	457

Notes. Probit regressions, marginal effects. Standard errors are in parentheses, clustered by matching group.

*** $p < 0.01$.

** $p < 0.05$.

* $p < 0.1$.

ones when the interaction partner is an honest seller. By contrast, tax-dishonest sellers treat e- and t-consumers more or less the same, as can be seen in Column (2).

Result 1. Undertreatment rates are lowest when both agents have opted for paying the tax in the first stage. Tax-dishonest consumers face significantly more undertreatment than tax-honest ones when interacting with an honest seller. Tax-dishonest sellers treat both types of consumers alike.

Overcharging does not result in any changes in efficiency since it is only a matter of redistribution. Table 2 shows that overcharging rates are very high across the board, which leaves little room for significant treatment differences in this fraud dimension. It is still worthwhile noting that overcharging is most frequent when two e-agents interact (Situation D), but the only statistically significant difference is between situations D and C ($p < 0.05$).

Result 2. We observe only small differences in overcharging rates between different types of interacting dyads. The highest overcharging rates occur in the interaction of two tax-dishonest agents.

We conclude this subsection with a word on actual tax evasion, i.e., the case in which both agents are tax-dishonest and the transaction goes untaxed. Pooling situations A to C and comparing them with D allows us to address the question whether behavior differs in conditions Γ_e and Γ_t – or, put differently, to evaluate the impact of actual tax evasion. There is no significant difference in undertreatment rates ($p > 0.16$) between taxed and untaxed markets, while there is a significant difference in overcharging rates ($p = 0.02$). The result for the undertreatment rate is perhaps not surprising since we have argued that attempted tax evasion is already enough to increase misbehavior by sellers compared to mutual tax honesty. When, in addition to attempted tax evasion, actual tax evasion occurs, undertreatment does not increase further because it is already at a high level and also because pairs of tax-dishonest agents tend to display slightly lower undertreatment rates than pairs of differently-minded agents. We will return to this point in Section 5.

¹² Such a difference in the moral cost is predicted, for instance, by the 'guilt-of-blame' variant of the guilt-aversion model of Battigalli and Dufwenberg (2007). For the present context this variant of the theory predicts that an expert cares about the customer's ex post inference regarding the extent to which he was willing to let her payoff expectations down. See Beck et al. (2013) for an application of the guilt aversion model to a credence goods context.

¹³ In order to control for the fact that observations within a matching group of eight participants are not independent of each other, all subsequent statistical p -values are based on two-sided Wilcoxon signed rank tests based on matching group averages across all twenty periods (unless stated otherwise).

¹⁴ In this and in all subsequent regressions, standard errors are clustered by matching group in order to control for interdependencies within each group.

4.2. The market game: interaction rates

Turning to the consumer's participation decision, a close look at the relevant column of Table 2 reveals that the interaction rate (IR) of 0.85 in Situation A is significantly higher than in all other situations ($p < 0.05$ for all pairwise comparisons). Furthermore, consumers differentiate between t-sellers and e-sellers and condition their interaction behavior on the tax decision of the seller. t-Consumers, in particular, are much more willing to enter the market if they are interacting with a t-seller than with an e-seller, increasing their IR from 0.64 to 0.85 ($p < 0.05$). e-Consumers also take their trading partner's tax decision into consideration, however here the effect of meeting a dishonest seller runs into the opposite direction. Comparing Situation B, in which e-consumers meet t-sellers, with Situation D, in which they meet e-sellers, reveals that e-consumers weakly significantly increase their interaction rate by 7 percentage points when being matched with an e-seller ($p < 0.1$).

These results on consumer behavior are qualitatively confirmed by regressions (3) and (4) of Table 3, which explain the interaction decision of consumers by means of an array of factors, including the tax decision of the seller. To receive more detailed information we ran separate regressions for t- and e-consumers. The dummy for meeting a tax-dishonest seller shows up with signs in line with the qualitative differences from Table 2 for both types of consumers, again suggesting that t-consumers decrease and e-consumers increase their interaction rates as a response to meeting an e-seller. This effect is significant only for tax-honest consumers and it has a large negative coefficient of -0.22 , mirroring the pronounced differences in interaction rates seen in Table 2. The other factors influencing the probability of interaction have the expected signs: Interaction probabilities decrease in p^h suggesting that both t-consumers and e-consumers anticipate that a higher price for the high-quality service increases the incentives for sellers to overcharge consumers. The effects of the price for the low-quality service are of opposite signs for t-consumers and e-consumers, and they are smaller in magnitude. The trend towards more selfish behavior of sellers indicated by Fig. 4 below results in a reduction of interaction frequencies over time, as indicated by the period-dummy in Table 2.

Result 3. Tax-honest consumers decrease their interaction rates significantly when they meet a tax-dishonest seller. Tax-dishonest consumers react in the opposite direction, but this effect is not significant.

4.3. The market game: overall efficiency

Combining the interaction and undertreatment rates in different situations to measure market efficiency, defined as $IR(1 - UR)$, we are able to investigate the efficiency loss that arises in the underlying market for credence goods as a consequence of attempted and actual tax evasion. As suggested by the results on undertreatment and interaction rates, efficiency is (with 0.48) highest in dyads with two tax-honest agents (Situation A in Table 2). This impression is confirmed by statistical tests: We find that efficiency in all other situations is significantly lower ($p < 0.05$) than in Situation A. All deviations from perfect honesty to attempted or successful tax evasion hence lead to a drop in efficiency of considerable magnitude, ranging from a drop by one third in Situation D to a drop by more than 50% in Situation C.

Result 4. Market efficiency is highest when both agents are tax-honest, but decreases significantly when at least one of the agents has revealed the intention to evade the tax. Attempted tax evasion is thus enough to lead to large efficiency losses.

Whenever tax evasion was attempted (and regardless of whether it was in fact realized through mutual agreement), the market suffered

from severe efficiency losses caused by a combination of an increase in the undertreatment rate and a decrease in the interaction frequency. Pooling situations A to C and comparing them with D reveals the efficiency impact of actual tax evasion. The similarity of overall undertreatment and interaction rates in market conditions Γ_e and Γ_t results in no significant difference in efficiency between the two market conditions ($p > 0.24$), although the difference is in the expected direction and is of quite substantial magnitude. Again, this is not really surprising since we have argued that attempted tax evasion is enough to reduce efficiency compared to mutual tax honesty.

4.4. The tax-decision stage: revealed intentions to evade taxes

Overall, in 60% of cases subjects were tax-dishonest, resulting in a total tax evasion rate of 37% (given that the consent of both trading partners is required for actual tax evasion). Tax (dis)honesty is a deliberate and consistent decision for most subjects: 84% of subjects can be classified as either being regularly tax honest (with tax-honesty rates of three quarters or more) or regularly tax-dishonest (with tax-honesty rates of one quarter or less). Fig. 3 reports the distribution of relative frequencies of subjects' tax decisions, broken down by sellers and consumers. A Kolmogorov–Smirnov test for equality in distributions shows no significant difference in tax-honesty of sellers and consumers ($p = 0.16$), even though, on average, sellers opt more often for tax evasion than consumers (63 vs. 57% of all decisions). At any rate, both distributions are clearly far from random ($p < 0.01$, binomial tests), suggesting that agents expect to interact in the market later on with positive probability. To get an impression of whether behavior is also stable over time, the black dashed line in Fig. 4 depicts the smoothed tax-dishonesty rate over all periods, and exhibits a slight upward trend.

To get more information on the drivers of behavior in the tax-decision stage of the game we run a probit regression as shown in Table 4. It explains revealed intentions to evade taxes in treatment ENDO by means of the participant type (seller or consumer), the two prices, the period, and idiosyncratic controls for gender, age, and risk attitudes. Contrary to the results of the non-parametric statistical test, the regression picks up a weakly significant difference between participant types, with sellers being about 12 percentage points more likely than buyers to opt for tax evasion, ceteris paribus. Similar to previous studies such as Baldry (1987) and Bayer and Sutter (2009) we also find a significant effect of gender on the probability of tax evasion, with females exhibiting a lower propensity to be tax-dishonest. Risk attitude shows up insignificantly, implying that we do not find any evidence that risk attitudes are important for tax-honesty in the tax-decision stage. Moreover, the prices p^h and p^l are insignificant.




Fig. 3. Tax-dishonest decisions per subject in ENDO. Note. Black (gray) bars indicate sellers (consumers).




Fig. 4. Behavior over time in treatment ENDO. Note. Smoothed by calculating 3-period moving averages.

5. What is driving the results?

In the previous section we have seen that market behavior in our main treatment differs remarkably between the four constellations of revealed intentions to evade taxes (i.e., between situations A–D in Table 2). However, within the *ENDO* treatment it was impossible to discriminate between the various effects that might be responsible for the differences.

For instance, we found that the undertreatment rate is only 43% in Situation A, while it is 15 percentage points higher in Situation B (with 58%). This finding is consistent with the existence of a *solidarity effect* which predicts that tax-honest sellers treat tax-honest customers better than tax-dishonest ones. However, the *punishment effect* makes exactly the same prediction regarding this comparison. This is not an accident – to the contrary, within the *ENDO* treatment the solidarity and conspiracy effect on the one hand, and the punishment effect on the other, are simply two sides of the same coin: Solidarity and conspiracy predict better performance in ‘symmetric situations’ (i.e., in A and in D), while punishment predicts worse performance in ‘asymmetric situations’ (i.e., in B and in C).

As a second example consider the finding that the undertreatment rate is significantly lower in A (where both parties are tax-honest) than in D (where both are tax-dishonest). This finding could be attributed to the solidarity effect (one would have to argue that it is stronger

Table 4
Tax decisions, treatment *ENDO*.

Dependent variable	Tax-dishonesty
Seller	0.122* (0.075)
p^l	0.001 (0.001)
p^h	–0.001 (0.001)
Period	0.002 (0.003)
Female	–0.214** (0.084)
Age	0.010 (0.011)
Risk factor	–0.021 (0.035)
Observations	1600

Notes. Probit regressions, marginal effects. Standard errors are in parentheses, clustered by matching group.

** $p < 0.05$.

* $p < 0.1$.

than the conspiracy effect), to the selection effect (which predicts that more pro-social sellers are not only less inclined to evade taxes but also less willing to defraud consumers), to the cake-size effect (one would have to argue that sellers behave in a more consumer-friendly way in an environment where earnings are intermediate with certainty than in the alternative scenario where gains from trade are high with high probability and low with low probability), or to any combination of those effects.

To receive information about the relative importance of the various effects for our findings we will exploit the two control treatments, *SELLER* and *EXO*. Panels (b) and (c) in Table 2 report the key variables for those treatments, and Table 5 shows regressions with undertreatment, overcharging and interaction rates as dependent variables and with the seven situations of Table 2 as independent variables (with Situation A as the left-out category).

First we search for evidence of a *cake size effect*. This effect could be responsible for differences between the untaxed Situation D and all other (taxed) situations in *ENDO*. The hypothesis would be that behavior differs between markets Γ_e and Γ_t due to the sheer change in profits and their probabilities, no matter how tax evasion has come about (in case it has). To check for the existence of a cake size effect we use our *EXO* treatment where – by design – the solidarity, conspiracy, punishment and selection effects cannot play any role, while the cake size effect is potentially present. Specifically, we compare Situation G to Situation H. We find only small differences in the provision and charging decisions of sellers across the two market situations. The overcharging rate is practically identical in G and H. The undertreatment rate increases slightly from 0.59 in taxed markets to 0.62 in untaxed markets. This difference is insignificant based on an F-test (Situation G = Situation H, post-estimation based on column 1 of Table 5) but it does show up as significant based on a Wilcoxon signed-ranks test ($p < 0.05$). Accordingly, we conclude by stating that we find some limited

Table 5
Undertreatment, overcharging and interaction rates, all treatments.

Dependent variable	(1)	(2)	(3)
	Undertreatment	Overcharging	Interaction
p^l	0.003 (0.002)	–0.004*** (0.001)	0.001 (0.002)
p^h	–0.007*** (0.002)	0.001 (0.001)	–0.008*** (0.002)
Period	0.009*** (0.003)	0.003*** (0.001)	–0.013*** (0.002)
Female	0.014 (0.066)	0.027 (0.024)	0.007 (0.028)
Age	–0.018 (0.012)	0.000 (0.002)	–0.009* (0.005)
Risk factor	–0.005 (0.028)	–0.000 (0.007)	0.018 (0.013)
Situation B	0.160* (0.082)	–0.007 (0.031)	–0.240*** (0.068)
Situation C	0.193* (0.086)	0.001 (0.038)	–0.238*** (0.027)
Situation D	0.120* (0.063)	0.044** (0.016)	–0.145** (0.071)
Situation E	0.148* (0.077)	–0.035 (0.053)	–0.193** (0.090)
Situation F	0.268*** (0.061)	0.038 (0.023)	–0.292*** (0.082)
Situation G	0.171*** (0.044)	0.006 (0.031)	–0.170** (0.083)
Situation H	0.176* (0.079)	–0.003 (0.035)	–0.174* (0.097)
Observations	1240	1962	2480

Notes. Probit regressions, marginal effects. Standard errors are in parentheses, clustered by matching group.

*** $p < 0.01$.

** $p < 0.05$.

* $p < 0.1$.

evidence for a cake size effect in the provision, but no evidence for a cake size effect in the charging decisions of sellers.

So far, we cannot disentangle the *selection effect* – i.e., the effect of self-selection by sellers into the underground economy – from the reaction of a seller to a consumer's tax decision, meaning that we cannot exclude the possibility that tax-evading sellers are per se more likely to undertake fraudulent behavior regardless of the tax attitude of consumers. To check for the presence of a selection effect (on top of the cake size effect) we use our *SELLER* treatment where by design the solidarity, conspiracy and punishment effects cannot play any role, while the selection and the cake size effect are potentially present. Specifically, we compare Situation E in Table 2 (where the seller has self-selected into the taxed economy) to Situation G (where the seller was randomly assigned to the taxed economy) and Situation F (where the seller has self-selected into the untaxed economy) to Situation H (where the seller was randomly assigned to the untaxed economy). Selection predicts higher undertreatment and overcharging rates in G than in E and in F than in H. Indeed, all four comparisons (two comparisons for each of the two fraud dimensions) are qualitatively in line with seller selection – undertreatment, for instance, is much higher among e-sellers in *SELLER* than under exogenously imposed tax evasion (70% vs. 62%). In terms of the relevant F-tests in Table 5, the differences are generally insignificant with the notable exception of significantly higher overcharging rates in situation F than in H (97% vs. 92%; $p < 0.05$, F-test based on Column 2 of Table 5).¹⁵ This indicates that those sellers who self-select into the untaxed economy may be more likely to overcharge their customers.

The *solidarity effect* and the *conspiracy effect* refer to the possibility that sellers may reserve better treatment for consumers of the same type. As we have defined those terms in the Introduction, solidarity may drive behavior in the relationship between t-sellers and t-consumers, while conspiracy might be relevant in the interaction between e-sellers and e-consumers, and both can alternatively be thought of as in-group favoritism. To check for the existence of a solidarity or a conspiracy effect we compare Situation A from treatment *ENDO* to Situation E from treatment *SELLER* and Situation D from treatment *ENDO* to Situation F from treatment *SELLER*. Solidarity between two tax-honest agents predicts that sellers cheat less in A than in E, because in A they are matched with a consumer of the same type, whereas in E nothing is known about the type of the consumer. For the same reason, conspiracy between two tax-dishonest agents predicts that undertreatment and overcharging will be less common in D than in F (where nothing is known about the consumer's type).

It is important to note that our identification strategy here is based on the implicit assumption that selection is the same across the *ENDO* and the *SELLER* treatment. This is not necessarily the case, as the implications of choosing to evade for a seller are different across the two conditions and, as such, could imply different selection processes. Specifically, a seller evading in *SELLER* realizes the benefits of evasion with certainty (he does not have to pay the taxes for sure), while this is not the case in *ENDO* (as here evasion takes place only if the buyer agrees). This limitation of our identification strategy should be kept in mind when interpreting our results.

While we have seen that overcharging frequencies are almost at a corner in every situation – and in this case higher in D than in F, contrary to the conspiracy hypothesis – the qualitative differences for undertreatment are fully in line with the existence of both a solidarity and a conspiracy effect. Undertreatment is less common in A than in E and in D than in F. Based on Column 1 of Table 5, both F-tests (Situation A = Situation E, respectively Situation D = Situation F) point to weakly significant differences ($p < 0.1$).

¹⁵ This difference, however, is not significant with non-parametric testing ($p > 0.1$, Mann-Whitney test based on matching group averages). All other differences between situations E and G, respectively F and H, are also insignificant based on Mann-Whitney tests, in line with the regression results reported here.

The *punishment effect* predicts that fraud is more prevalent in situations where the seller knows he is matched with someone of a different tax attitude than in situations where he does not have this information. To check for the existence of a punishment effect we compare Situation B from treatment *ENDO* to Situation E from treatment *SELLER* and Situation C from treatment *ENDO* to Situation F from treatment *SELLER*. Here we exploit the fact that the cake size effect was found to be of minor importance and we assume again that selection is the same across the *ENDO* and the *SELLER* treatment (which is not necessarily the case – as indicated above). Punishment then predicts more consumer-friendly behavior in E than in B (where the seller knows that the consumer is of a different type) and more consumer-friendly behavior in F than in C (for the same reason). The corresponding differences in undertreatment and overcharging rates are either very small or in the “wrong” direction (such as the comparison between the 65% undertreatment rate in C and the 70% UR in F) and the F-tests in Table 5 show that the differences are always insignificant. Hence, we find no evidence to support the presence of a punishment effect in our experiment. We summarize the findings of this section as:

Result 5. We find evidence for the existence of a solidarity effect between two tax-honest agents and a conspiracy effect between two tax-dishonest agents, leading to less fraud when agents reveal the same tax attitude. We also find some evidence for the existence of a selection effect, with tax-dishonest sellers being more likely to overcharge their customers, but almost no evidence for the existence of a punishment effect or a cake size effect.

In this section's analysis we have so far mainly dealt with the various effects driving the provision and charging behavior of sellers, while remaining silent about the behavior of consumers and the resulting interaction rates. The reason is that our design does not allow us to cleanly identify the relative roles of the many effects that might be driving the decision to enter the market or not. In particular, entry decisions by consumers can be driven by a solidarity, a conspiracy, a punishment or a cake-size effect, but also by the consumer's expectation of (some of) these effects affecting their partners as sellers. To give an example of this, let us compare the interaction rates in situations A and C. The fact that interaction is much more frequent in A than in C may be driven by a solidarity effect (t-consumers reward t-sellers' tax-honesty in A by increasing their propensity to enter the market), a punishment effect (t-consumers punish e-sellers' tax-dishonesty in C by decreasing their propensity to enter), or an anticipation of a number of different things: Anticipation of less (more) fraud by the seller in A (in C) due to the solidarity (punishment) effect on the seller's side, or anticipation of more fraud in C due to the perceived selection effect – in this latter case, the consumer perceives the seller's tax evasion decision as a signal of his preferences and of the resulting provision and charging behavior. Similarly, the higher interaction rate in D than in B may be driven by the conspiracy effect, the punishment effect, or by the consumer's anticipation (e.g., if consumers hope for better treatment in D due to conspiracy from the seller).

Our controls are of limited help in disentangling those effects. For instance, following the line of argument that we used for undertreatment, solidarity predicts higher interaction rates in A than in E and conspiracy predicts higher interaction in D than in F. We do find significant differences in both of these comparisons ($p < 0.05$; F-tests based on column 3 of Table 5), which is consistent with the existence of a solidarity and a conspiracy effect in the consumers' interaction decisions. At the same time, however, consumers may interact more frequently in A and in D because they expect the seller to be driven by the solidarity or the conspiracy effect in these situations. This issue could potentially be addressed by eliciting consumers' beliefs about the behavior of the seller or by running a number of additional control treatments. Given that the main focus of our paper is on the interplay between revealed intentions to pay or evade taxes and provision and charging behavior in

the underlying credence goods market, we leave this an open question for future research.

6. Conclusion

The channels through which tax evasion can be detrimental to welfare are intricate. We have experimentally investigated the effects of revealed intentions to pay or evade taxes on the performance of a particular kind of market that is prone to inefficiencies in trade. Participants in a credence goods market were given the opportunity to reveal their intentions regarding tax evasion before the actual transaction took place. Only when both the seller and the consumer opted for tax evasion was the transaction actually concealed from the tax authority.

We are contributing to the literature on tax evasion by demonstrating that the phenomenon is detrimental not only because it yields an unfair distribution of the overall tax burden, wastes resources when taxpayers try to conceal and tax authorities try to detect tax evasion, imposes uncertainty on risk-averse evaders and distorts competition, but also because of its negative impact on the efficiency of interaction in the underlying market. In transactions in which at least one partner was tax-dishonest, efficiency was up to 50% lower than in a constellation with exclusively tax-honest agents. This difference is driven by higher undertreatment frequencies and lower interaction rates in markets in which tax evasion was attempted or successfully conducted. We found that the efficiency drop in untaxed markets is hardly driven by the change in market conditions. Rather, efficiency losses arise mainly due to the fact that subjects who prefer to pay taxes react negatively to the intent of a trading partner to evade them by an increased frequency of undertreatment (in the case of sellers) and a decreased frequency of interaction (in the case of consumers). This is in line with our initial conjecture about the existence of a solidarity effect between two tax-honest agents that seems to reduce the undertreatment rate and increase the interaction rate in pairs of tax-honest market participants. Although we also find evidence for a conspiracy effect between two tax-dishonest agents, which *ceteris paribus* would lead to more consumer-friendly behavior in the underground sector, consumers do not really profit from that effect because it is almost entirely canceled out by a selection effect that causes more selfish sellers to self-select into the underground sector.

While our work provides a first indication of the presence of hidden costs of tax evasion, the empirical importance of these costs should be substantiated in future lab and field experiments, perhaps under variations in the institutional setup. Our findings so far suggest that consumers are well advised to be cautious of proposing to evade taxes in situations in which a seller has an informational advantage over the consumer. If the seller is tax-honest then he might treat the consumer worse because he has less scruple to defraud a tax-dishonest than a tax-honest consumer, due to the fact that he does not identify himself with a tax-dishonest person. In this case, proposing tax evasion leads to deteriorated service without any compensating benefit (given that tax evasion requires mutual consent). In the alternative case where the seller agrees to conceal the transaction, the myopic gain of the evaded tax may easily be canceled out by the deteriorated service provision resulting from the fact that a tax-dishonest seller is more likely to commit fraud in the first place. This latter statement is supported by our finding that selfish sellers tend to self-select into the shadow economy. So, before agreeing to a transaction, it seems advisable for a consumer to try to elicit the tax attitude of the seller who has the power to commit fraud and to abstain from a purchase if she discovers on the seller's side an intention to evade taxes.

Acknowledgements

We are indebted to the co-editor, Henrik Kleven, and two anonymous referees for their excellent comments which helped to improve the paper a lot. We also wish to thank Kai Konrad and participants at

the X. Tax Day of the Max Planck Institute for Tax Law and Public Finance for helpful comments. Financial support from the Austrian Science Fund (FWF) through grant numbers P20796 and P22669 is gratefully acknowledged.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <http://dx.doi.org/10.1016/j.jpubeco.2015.06.003>.

References

- Allingham, M.G., Sandmo, A., 1972. Income tax evasion: a theoretical analysis. *J. Public Econ.* 1, 323–338.
- Alm, J., McClelland, G., Schulze, W., 1992. Why do people pay taxes? *J. Public Econ.* 48, 21–38.
- Andreoni, J., Erard, B., Feinstein, J., 1998. Tax compliance. *J. Econ. Lit.* 36, 818–860.
- Baldry, J.C., 1986. Tax evasion is not a gamble: a report on two experiments. *Econ. Lett.* 22, 333–335.
- Baldry, J.C., 1987. Income tax evasion and the tax schedule: some experimental results. *Public Financ.* 42, 357–383.
- Battigalli, P., Dufwenberg, M., 2007. Guilt in games. *Am. Econ. Rev.* 97, 170–176.
- Bayer, R.-C., 2006. A contest with the taxman — the impact of tax rates on tax evasion and wastefully invested resources. *Eur. Econ. Rev.* 50, 1071–1104.
- Bayer, R.-C., Sutter, M., 2009. The excess burden of tax evasion — an experimental detection-concealment contest. *Eur. Econ. Rev.* 53, 527–543.
- 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.
- Boadway, R., Marceau, N., Mongrain, S., 2002. Joint tax evasion. *Can. J. Econ.* 35, 417–435.
- Brandts, J., Charney, G., 2011. The strategy versus the direct-response method: a first survey of experimental comparisons. *Exp. Econ.* 14, 375–398.
- Casari, M., Cason, T., 2009. The strategy method lowers measured trustworthy behavior. *Econ. Lett.* 103, 157–159.
- Chang, J., Lai, C., 2004. Collaborative tax evasion and social norms: why deterrence does not work. *Oxf. Econ. Pap.* 56, 344–368.
- Charney, G., Rigotti, L., Rustichini, A., 2007. Individual behavior and group membership. *Am. Econ. Rev.* 97, 1340–1352.
- Chetty, R., 2009. Is the taxable income elasticity sufficient to calculate deadweight loss? The implications of evasion and avoidance. *Am. Econ. J.: Econ. Policy* 1, 31–52.
- Clotfelter, C., 1983. Tax evasion and tax rates: an analysis of individual returns. *Rev. Econ. Stat.* 65, 363–373.
- Dulleck, U., Kerschbamer, R., 2006. On doctors, mechanics, and computer specialists: the economics of credence goods. *J. Econ. Lit.* 44, 5–42.
- Dulleck, U., Kerschbamer, R., Sutter, M., 2011. The economics of credence goods: on the role of liability, verifiability, reputation and competition. *Am. Econ. Rev.* 101 (2), 526–555.
- Erard, B., Feinstein, J.S., 1994. The role of moral sentiments and audit perceptions in tax compliance. *Public Financ.* 49, 70–89.
- Feinstein, J.S., 1991. An econometric analysis of income tax evasion and its detection. *RAND J. Econ.* 22, 14–35.
- Feldstein, M., 1995. The effect of marginal tax rates on taxable income: a panel study of the 1986 tax reform act. *J. Polit. Econ.* 103, 551–572.
- Feldstein, M., 1999. Tax avoidance and the deadweight loss of the income tax. *Rev. Econ. Stat.* 81, 674–680.
- Fischbacher, U., 2007. z-Tree: Zurich toolbox for ready-made economic experiments. *Exp. Econ.* 10, 171–178.
- Frey, B., Torgler, B., 2007. Tax morale and conditional cooperation. *J. Comp. Econ.* 35, 136–159.
- Gordon, J.P.P., 1989. Individual morality and reputation costs as deterrents to tax evasion. *Eur. Econ. Rev.* 33, 797–805.
- Götte, L., Huffman, D., Meier, S., Sutter, M., 2012. Competition between organizational groups: its impact on altruistic and anti-social motivations. *Manag. Sci.* 58, 948–960.
- Greiner, B., 2004. An online recruiting system for economic experiments. In: Kremer, Kurt, Macho, Volker (Eds.), *Forschung und wissenschaftliches Rechnen 2003. Gesellschaft fuer wissenschaftliche Datenverarbeitung, GWDG Bericht 63*, Goettingen, pp. 79–93.
- Güth, W., Strauß, S., Sutter, M., 2005. Tax evasion and state productivity — an experimental study. *Metroeconomica* 56, 85–100.
- Holt, C.A., Laury, S.K., 2002. Risk Aversion and Incentive Effects. *American Economic Review* 92 (5), 1644–1655.
- Huck, S., Lünser, G., Tyran, J.-R., 2010. Consumer networks and firm reputation: a first experimental investigation. *Econ. Lett.* 108, 242–244.
- Huck, S., Lünser, G., Tyran, J.-R., 2012. Competition fosters trust. *Games Econ. Behav.* 76, 195–209.
- Huck, S., Lünser, G., Tyran, J.-R., 2013. Price competition and reputation in markets for experience goods: an experimental study. Discussion Paper No. SP II 2013-312, Wissenschaftszentrum Berlin für Sozialforschung, Abteilung Ökonomik des Wandels.
- IRS, 2011. The tax gap www.irs.gov/newsroom/article/0,,id=158619,00.html, U.S. Department of the Treasury, Internal Revenue Service (Accessed June 2015).

- IRS, 2012. IRS Releases 2006 Tax Gap Estimates. www.irs.gov/uac/IRS-Releases-2006-Tax-Gap-Estimates, U.S. Department of Treasury Internal Revenue Service (Accessed June 2015).
- Kerschbamer, R., Sutter, M., Dulleck, U., 2015. How social preferences shape incentives in (experimental) markets for credence goods. *Econ. J.* (forthcoming).
- Kleven, H.J., Kreiner, C.T., Saez, E., 2009. Why can modern governments tax so much? An agency model of firms as fiscal intermediaries. *NBER Working Paper* 15218.
- Kleven, H.J., Knudsen, M.B., Kreiner, C.T., Pedersen, S., Saez, E., 2011. Unwilling or unable to cheat? Evidence from a tax audit experiment in Denmark. *Econometrica* 79, 651–692.
- Konrad, K., Qari, S., 2012. The last refuge of a scoundrel? Patriotism and tax compliance. *Economica* 79, 516–533.
- Madzharova, B., 2004. The Effect of Low Corporate Tax Rate on Payroll Tax Evasion. In: Fuest, C., Zodrow, G.R. (Eds.), *Critical Issues in Taxation and Development*. MIT Press, Cambridge, MA.
- Naritomi, J., 2013. Consumers as Tax Auditors. *Mimeo*. Harvard University.
- Reinganum, J., Wilde, L., 1985. Income tax compliance in a principal-agent framework. *J. Public Econ.* 26, 1–18.
- Saez, E., Slemrod, J., Giertz, S.H., 2012. The elasticity of taxable income with respect to marginal tax rates: a critical review. *J. Econ. Lit.* 50, 3–50.
- Schneider, F., Enste, D.H., 2000. Shadow economies: size, causes, and consequences. *J. Econ. Lit.* 38, 77–114.
- Slemrod, J., 1998. Methodological issues in measuring and interpreting taxable income elasticities. *Natl. Tax J.* 51, 773–788.
- Slemrod, J., 2007. Cheating ourselves: the economics of tax evasion. *J. Econ. Perspect.* 21, 25–48.
- Spicer, M., Becker, L., 1980. Fiscal inequity and tax evasion: an experimental approach. *Natl. Tax J.* 33, 171–175.
- Strand, J., 2005. Tax distortions, household production, and black-market work. *Eur. J. Polit. Econ.* 21, 851–871.
- Torgler, B., 2002. Speaking to theorists and searching for facts: tax morale and tax compliance in experiments. *J. Econ. Surv.* 16, 656–683.
- Torgler, B., 2004. Moral suasion: an alternative tax policy strategy? Evidence from a controlled field experiment in Switzerland. *Econ. Gov.* 5, 235–253.
- Yaniv, G., 1992. Collaborated employee–employer tax evasion. *Public Financ* 47, 312–321.
- Yitzhaki, S., 1979. A note on optimal taxation and administrative costs. *Am. Econ. Rev.* 69, 475–480.
- Yitzhaki, S., 1987. On the excess burden of tax evasion. *Public Financ. Q.* 15, 123–137.