THE ECONOMIC JOURNAL

The Economic Journal, Doi: 10.1111/ecoj.12260 © 2015 Royal Economic Society. Published by John Wiley & Sons, 9600 Garsington Road, Oxford OX4 2DQ, UK and 350 Main Street, Malden, MA 02148, USA.

SECOND-DEGREE MORAL HAZARD IN A REAL-WORLD CREDENCE GOODS MARKET*

Loukas Balafoutas, Rudolf Kerschbamer and Matthias Sutter

In a field experiment in the market for taxi rides we investigate a phenomenon called second-degree moral hazard – the tendency of the supply side in a market to react to anticipated moral hazard on the demand side by increasing the extent or price of the service. Our moral hazard manipulation consists of some passengers explicitly stating that their expenses will be reimbursed. This has a strong positive effect on the likelihood and the amount of overcharging and consequently increases consumer expenditure. Our results suggest that second-degree moral hazard may have a severe impact on the provision of credence goods.

Studying the crucial role of asymmetric information on market outcomes has a long tradition in economic theory, in particular in models that demonstrate market failure when asymmetrically informed agents interact with each other. A frequently investigated phenomenon is moral hazard (Pauly, 1974; Holmström, 1979), which refers to the tendency of agents who do not bear full responsibility for the costs or risks associated with their actions to exercise less effort to reduce these costs. When applied to markets for insurance, the moral hazard hypothesis states that agents with a more comprehensive insurance coverage have weaker incentives to avoid exposure to risk, meaning that insured events will more often occur with fully insured agents than among those with imperfect or no coverage. A number of empirical studies have tested for the presence of moral hazard in markets plagued by asymmetric information, with generally mixed results. For instance, Chiappori *et al.* (1998) and Einav *et al.* (2013) find some evidence of moral hazard in the demand for medical services for insured agents, while Chiappori and Salanié (2000) and Doran *et al.* (2005) reject the importance of moral hazard in the markets for automobile insurance and for prescription medicine respectively.

In an organisational context, moral hazard can manifest itself in the relation between owners and managers of a firm in the case of misaligned incentives and unobservable employee effort (Grossman and Hart, 1983; Mookherjee, 1984; Bartling and von Siemens, 2010). One example of employee effort that is subject to moral hazard is the investment in cost-reduction regarding expense accounts, travel costs and company cars or, more generally, any 'hidden' effort geared towards cost minimisa-

* Corresponding author: Loukas Balafoutas, Department of Public Finance, University of Innsbruck, Universitaetsstrasse 15, A-6020 Innsbruck, Austria. Email: loukas.balafoutas@uibk.ac.at.

We thank Marco Castillo, Andrea Galeotti, Nikos Nikiforakis, Ragan Petrie, Achim Zeileis, two anonymous referees, as well as audiences at the Universities of Innsbruck, Cologne and the Georgetown University School of Foreign Service in Qatar for very helpful comments. Special thanks are due to Konstantinos Konstantakis, Vasileios Selamis, Christina Tseligka, Amaryllis Zeppou and Nancy Karystinou for assisting in the experiment. Financial support from the Austrian Science Fund (FWF grants P20796 and P26901) and from the Australian Research Council (ARC Grant DP110103653) is gratefully acknowledged.

tion. The presence of moral hazard in this particular context is an important consideration among practitioners in the field.[1]

The existing empirical studies on moral hazard focus exclusively on the direct impact of asymmetric information on market outcomes, thereby ignoring possible repercussion effects. As an illustration consider the market for health care services and assume that the consumer of the service – in this case, a patient – is fully insured and interacts with a seller of the service – in this case, a physician. Moral hazard implies that the patient may have incentives to demand more of the service than required (by asking for more numerous or more extensive tests or treatments), since he will not bear its costs. However, the behaviour of the physician may also be affected by the extent of the coverage: if the physician expects that the patient is not concerned about minimising costs, he may be more inclined to suggest or prescribe more expensive treatments. Notice that the two stories – which we will call first-degree moral hazard and second-degree moral hazard – are observationally equivalent in terms of final outcomes, in the sense that more extensive insurance coverage leads to higher expenditure, but the mechanisms are different. While first-degree moral hazard operates through the demand side, second-degree moral hazard increases expenditure through supplier-induced demand – the artificial increase in demand induced by the actions of the seller.

A similar second-degree moral hazard story also applies to the above example on the lack of cost-minimising incentives by firm employees. Consider the case of an employee who must travel for work purposes and is responsible for booking the tickets himself. Since he will be fully reimbursed for his expenses, his incentives for finding and buying the lowest fare are weak (first-degree moral hazard). Moreover, if he visits a travel agent and informs him of the purpose of travel, the agent may deliberately suggest more expensive travel options or charge a higher-than-justified commission, anticipating that the employee will not put any effort into monitoring or protesting the agent's actions (second-degree moral hazard).

In this article, we present the findings from a field experiment that studies the impact of full reimbursement (or coverage) of expenses on *ex post* realised expenditures induced by the actions of the service provider – and not of the consumer. To our knowledge, this is the first study to examine the possibility that second-degree moral hazard may be partly responsible for the positive correlation between coverage and realised expenditure. A further contribution of our article is that, to the best of our knowledge, it is the first to use a controlled field experiment to study (second-degree) moral hazard and its impact on market outcomes (for a general discussion on field experiments and the advantages of this methodology, see List (2006) and List and Reiley (2008)).

We have conducted our study in the market for taxi (cab) rides. In the case of taxi rides in an unknown city, the service traded on the market is a credence good (Darby and Karni, 1973), meaning that an expert seller possesses superior information about

[1] See, for example a CNN report on expense account fraud and the extent to which business travellers tend to inflate their expenses (including taxi receipts): http://edition.cnn.com/2011/12/05/travel/expense-account-business-travel/index.html (accessed: 19 September 2014).

the needs of the consumer.[2] In particular, the driver knows the correct route to a destination while the consumer does not. This property of credence goods opens the door to different types of fraud: overtreatment occurs when the consumer receives more extensive treatment than what is necessary to meet his needs (with taxi rides, this amounts to a time-consuming detour); in the opposite case of undertreatment, the service provided is not enough to satisfy the consumer (i.e. he does not reach his destination); finally, in credence goods markets where the consumer is unable to observe the quality she has received, there might also be an overcharging incentive, meaning that the price charged by the seller is too high, given the service that has been provided.

In the context of credence goods markets such as taxi rides or medical and repair services, second-degree moral hazard seems particularly relevant because the informational asymmetries present in such markets facilitate supplier-induced demand, or even fraud; see Sülzle and Wambach (2005) for a theoretical model. Anecdotal evidence suggests that in such markets second-degree moral hazard can be economically even more important than first-degree moral hazard. For example, a report in the German weekly magazine *Der Spiegel* from 23 December 2012 estimates the yearly damage to insurers from faked and inflated medical bills amounts to 6–24 billion euro.[3] Inflated bills seem to be more of a supply than a demand-side phenomenon. As we argue in the conclusion, in some credence goods sectors the phenomenon seems to be so important that it has led to the design of fraud-restraining institutions.

In a recent field experiment, Balafoutas *et al.* (2013) have examined the role of informational asymmetries between sellers and buyers in the market for taxi rides and established that fraud by taxi drivers is very systematic in the sense that it increases with the extent of the informational asymmetry between the driver and the passenger. In their paper, all treatment variations were restricted to the extent to which drivers were better informed than passengers. No third party was involved, making it impossible to investigate how second-degree moral hazard affects the supply side of the market. This is what we do in this article. Basically, we use the same approach as in Balafoutas *et al.* (2013) but keep the informational asymmetry between the driver and the passenger constant and introduce a third party into the interaction, namely a (fictional) employer who reimburses the passenger's costs.[4,5]

[2] For some theoretical and experimental work on the properties of credence goods markets see Dulleck and Kerschbamer (2006), and Dulleck *et al.* (2011). Related work on the effects of informational asymmetries in experience goods markets is by Huck *et al.* (2007, 2012).

[3] See http://www.spiegel.de/wirtschaft/krankenkassen-detektive-jagen-betruegerische-aerzte-mit-spezial-software-a-873059.html

[4] There is an interesting literature on the labour supply of (New York City) taxi drivers (Camerer *et al.*, 1997; Farber, 2005, 2008; Crawford and Meng, 2011). However, this literature remains silent on the question of whether and to what extent taxi drivers exploit their informational advantage over customers in the provision of this credence good. Hence, we do not go into the details of these studies.

[5] Schneider (2010) and Jackson and Schneider (2011) use data from New York City taxi drivers in order to estimate the extent and determinants of moral hazard in a different dimension, namely in the relationship between owners and lessees of vehicles. As we explain in subsection 1.2, this aspect is not relevant in the environment of our study.

Specifically, in our study a team of research assistants took 400 undercover taxi rides in the capital city of Greece, Athens. The assistants, acting as passengers on 11 different routes throughout the city, always revealed to the driver that they were unfamiliar with the city, thus giving the taxi ride the characteristic of a typical credence good. Our main experimental variation consisted solely of a short phrase in one of the two treatments, in which the passengers indicated that they needed a receipt in order to have their expenses reimbursed by their employer. We call this the 'moral hazard treatment', in the sense that the passenger conveyed the impression of not personally incurring the costs of the fare and was therefore arguably perceived by the driver as less likely to notice, or report, fraudulent behaviour on his side.

We find that our moral hazard manipulation has an economically pronounced and statistically significant positive effect on the likelihood and the amount of overcharging, with passengers in that treatment being about 17% more likely to pay higher-than-justified prices for a given ride. This also leads to significantly higher consumer expenditures in this treatment on average. At the same time, the rate of overtreatment (by taking time-consuming detours) does not differ across treatments. Hence, second-degree moral hazard does not increase the extent of overtreatment compared to the control, while it does increase the likelihood and the extent of overcharging.

Before we proceed to describe in detail the experiment and its findings, we stress once more that first-degree or direct moral hazard (on the consumer side) is ruled out in this experiment since the behaviour of the passengers is exogenously controlled by the experimenters and kept constant across treatments. Our research question is how the perception that moral hazard exists in the relationship between the consumer and the reimbursing employer affects the seller of the service. We have called this mechanism second-degree moral hazard, and whenever the term moral hazard shows up in the following it refers to this kind of mechanism. Moreover, we are not addressing potential effects of adverse selection (Akerlof, 1970; Rothschild and Stiglitz, 1976). Generally speaking, a challenge in the empirical literature on asymmetric information and its impact on market outcomes is to distinguish between adverse selection and moral hazard, since both predict a positive correlation between coverage and *ex post* realisations of risk, albeit through very distinct mechanisms.[6] Adverse selection implies that high-risk agents *ex ante* self-select into contracts with more extensive coverage, while moral hazard means that agents may increase their risk *ex post* because they possess more extensive coverage. Our experimental methodology rules out adverse selection because agents (passengers) do not choose between the two treatments, but are randomly assigned to one of them. This allows us to sidestep the problem of disentangling moral hazard from adverse selection, which we view as a further advantage of our study.

[6] For instance, Chiappori and Salanié (2000, p. 60) acknowledge that 'the general problem of distinguishing between adverse selection and moral hazard from insurance data is left for future research'. A recent attempt to disentangle the two phenomena using automobile insurance data from France is Dionne *et al.* (2013).

1. Experimental Design and Implementation

1.1. *Experimental Methods*

The field experiment was run by means of taking undercover taxi rides. Four research assistants (two males and two females) presented themselves as customers and documented the driving and charging decisions of taxi drivers (who were unaware that their behaviour was being studied).[7] Each observation consisted of an assistant entering a taxi and requesting to be taken to a particular destination, following a fixed script: 'I would like to get to [name of destination]. Do you know where it is? I am not from Athens'. This is the script used in Balafoutas *et al.* (2013) for the role of non-local native passengers – that is, for passengers who speak the native language but clearly indicate to the driver that they are not familiar with the city. Arguably, this means that the taxi ride is considered a credence good by the driver, in the sense that the driver expects to have an informational advantage that he or she may try to exploit.[8]

We implemented two treatments. In the control treatment (henceforth CTR), a few seconds after the ride had begun the passenger added the following question to the driver: 'Can I get a receipt at the end of the ride?' In the moral hazard treatment (henceforth MOH) the same question was supplemented by a short phrase explaining that the passenger would have his or her expenses reimbursed: 'Can I get a receipt at the end of the ride? I need it in order to have my expenses reimbursed by my employer'. The two treatments were exactly identical except for this brief additional phrase, allowing us to identify the potential effects of this revealed information. Our hypothesis is that fraudulent behaviour will be more prevalent or more pronounced in the moral hazard treatment, since – as explained in the introduction – drivers may infer that passengers in this treatment have weak incentives to control, or to report, overtreatment or overcharging.[9] Even though our hypothesis does not differentiate between the two fraud dimensions (overtreatment and overcharging) and predicts that both will be more widespread or more pronounced in treatment MOH, overcharging is the more attractive and hence more likely dimension, as we will outline in more detail in subsection 3.1.

In addition to this treatment manipulation, we investigated gender effects in service provision and the potential interaction of gender and moral hazard by recruiting two male and two female research assistants. Even though we had no formal prior hypothesis regarding the role of gender in the context of this experiment, intuition

[7] All research assistants were Greek natives from Athens who were hired and paid a fixed compensation for their involvement in the experiment. They were instructed in detail about the scenarios that they had to enact and about how to collect and document the data.

[8] This is in sharp contrast to the field study by Castillo *et al.* (2013) on gender differences in bargaining outcomes in the market for taxicab rides in Lima, Peru. In that study the undercover customers are locals who know the shortest route to their destination. This implies that the core credence good problem in the market for taxi rides (overtreatment in the form of taking circuitous routes) is absent. Also, taxis in Lima do not have posted prices and meters – prices are rather set by face-to-face negotiations. This implies that overcharging is not an issue in the Lima study either.

[9] Note that transportation authorities have, in principle, the power to confiscate the taxi licence in case of fraud.

Table 1
Number of Observations by Treatment and Gender

Treatment	CTR	MOH	Total
Male passengers	100	100	200
Female passengers	100	100	200
Total	200	200	400

might lead one to expect that women fall victim to fraud more often than men, for instance, because they are perceived as less likely to engage in confrontation with the driver. Table 1 summarises our design (including the number of collected observations).

Rides were always taken in quadruples, meaning that all four assistants took taxis from the same origin to the same destination, in intervals of one to two minutes between each other. This design feature was implemented in order to control for factors unrelated to the treatment manipulation but possibly related to the optimal route and to the price of the ride, such as traffic and weather conditions, unforeseen or time-specific events such as closed roads etc. Within each quadruple, the order of entering the taxi was random, and two passengers (one female and one male) enacted the control treatment, while the other two enacted the moral hazard treatment. As a result, all four cells of the experimental design (Table 1) are represented in each quadruple.

1.2. *The Market Environment*

The market for taxi rides is regulated nationwide in Greece, both in terms of market entry and details of the tariff system. A government authority issues taxi licences as perpetuities for their holders. Currently, approximately 14,000 taxi licences are valid in Athens, implying a ratio of roughly 350 taxis per 100,000 inhabitants. This is a considerably larger supply than, for instance, in London with 280 taxis, Berlin with 210 taxis, New York City with 160 taxis (yellow cabs only) or the whole US with around 110 taxis per 100,000 inhabitants.[10]

The tariff system is regulated such that there is a fixed fee of €1.19 and a variable part. The variable part is either distance-dependent – in this case the tariff is €0.68 per kilometre during daytime (i.e. from 5 AM until midnight) and €1.19 per kilometre during nighttime – or duration-dependent – then it yields €0.1808 per minute both during daytime and during nighttime. The algorithm for charging is standardised nationwide in all taximeters and switches automatically to the counting method (distance-dependent *versus* duration-dependent) that is more profitable for the driver. Given the large supply of taxis in Athens, drivers typically have to queue for passengers for long periods of time. This implies that it is generally far more

[10] Source: Own calculations, based on numbers from Schaller Consulting (2006), 'Transportation for London' (http://www.tfl.gov.uk/businessandpartners/taxisandprivatehire/1380.aspx), 'Taxi Innung Berlin' (http://www.taxiinnung.org/Taxi-Bestellen.24.0.html; accessed: 19 September 2014), and Schaller (2005).

profitable for a taxi driver to take a passenger on a detour – thus providing overtreatment – than to choose the shortest and quickest route in the hope to accumulate many fixed fees. Moreover, abstracting from possible punishment, there are clear incentives for overcharging, since it increases the driver's revenue without affecting the cost of service. Note that the marginal incentives for taxi drivers to engage in fraud of any type are practically identical for owner-drivers and for drivers who lease the vehicle, because leasing a taxi comes at a fixed cost (of roughly €35 per shift). Therefore, taxi drivers are always residual claimants, meaning that the possible profits from fraud are reaped by themselves.

Since January 2013, every consumer in Greece has the right to refuse payment for a transaction of a good or service if a legal receipt has not been issued. Accordingly, every taxi is equipped with a small cash register, which is connected to the taximeter and automatically produces a receipt at the end of each ride. This implies that asking for a receipt should have no effect *per se* and our assistants only did so to have an anchor for adding the phrase explaining that they are having their expenses reimbursed (in the moral hazard treatment) and to keep the script identical across treatments as much as possible.

1.3. *Data Set*

The experiment was conducted over 15 days in March 2013 and July 2014, covering each day of the week at least once. All rides were taken between 8 AM and midnight on 11 different routes throughout the city of Athens, randomising among routes during the day (meaning that data on each route were collected during several days and that several routes were driven each day, in random order).[11] We collected a total of 400 observations, organised in 100 quadruples of (almost) simultaneous rides.[12] The mean length of a ride was 16.2 km and the mean duration was 23.4 minutes, thus amounting to a total of about 6,482 kilometres or 156 hours of driving.

For each observation, our assistants collected the following data: date and time at the start of the ride; route; duration; distance driven; total price (fare) paid; gender and approximate age of the driver; vehicle manufacturer; weather and traffic conditions. The exact distance driven was measured with a portable GPS satellite logger (see Figure A1 in Appendix A) that records the chosen route and the taxi's exact position and speed at each point in time. With these data, we could quantify overtreatment in the form of detours and also determine the correct fare for the driven distance. This helped us to create an overcharging indicator, which is one if a driver asked for more than justified by the driven distance, and zero otherwise.

In order to classify the source of overcharging, we resorted to the detailed information collected by our assistants, describing whether the driver artificially increased the fare for a given ride (e.g. by demanding unjustified extras or applying an

[11] For a list of all routes and brief description of the points of origin and destination, please refer to Tables A1 and A2 in Appendix A.

[12] We note that the sample initially consisted of 256 observations collected in 2013. The remaining 144 observations were collected in 2014 at the advice of the editor and referees and resulted in stronger statistical significance. Based on the initial sample alone, some of the results outlined in subsection 2.2 were only marginally significant.

incorrect tariff). Even though any such actions were taken by the driver in secrecy, our assistants were thoroughly instructed about the correct price list and about possible 'tricks' employed by the drivers, so that they were able to identify the method of fraud. We also computed the overcharging amount, defined as the difference between the total price and the price that should have been paid for the actually driven distance – or, in other words, the amount by which the customer was cheated (on top of potential overtreatment through detours). We classify an observation as a case of overcharging if the amount of this markup is at least 5% of the total price paid, noting that all our results are robust to extending the definition of overcharging to include lower amounts.[13]

2. Results

2.1. *Overtreatment*

In order to quantify overtreatment we have constructed a duration index by dividing the duration of each ride by the minimum duration recorded in that particular quadruple. The index thus controls for all unobserved characteristics related to driving conditions (such as traffic). The duration of a ride is, however, not the sole dimension that relates to overtreatment: one must also take the distance driven into account, which will be larger if a driver takes detours in order to increase the total fare. Hence, we also construct a distance index, defined as the ratio between the distance in each observation and the minimum distance in that particular quadruple.

The values of the two overtreatment indices (duration and distance index) by treatment and gender are shown in Table 2. Panel (*a*) reveals that there are no substantial differences in overtreatment across treatments or across gender, since the duration index is practically identical across treatments (1.14 in CTR and 1.13 in MOH) and across gender (1.13 for males and 1.14 for females). This impression is confirmed by statistical tests. In particular, we test for treatment differences by taking the mean value of the duration index for the two passengers in CTR in a given quadruple and comparing it with the mean value of the two passengers in MOH in the same quadruple, in order to account for the fact that observations are not independent within each quadruple due to the way the index is constructed. This leads to a total of 100 observations on paired samples and the null hypothesis of no significant treatment difference is not rejected by Wilcoxon signed-ranks tests ($p = 0.44$).[14] Moreover, men and women are also not taken on rides of significantly different length, on average. Panel (*b*) reports the values of our second measure of overtreatment, namely the distance index, showing that the value of this index was slightly higher in the moral hazard treatment. This difference is, however, only marginally significant ($p = 0.088$, Wilcoxon signed-ranks test). Hence, we are led to conclude that our treatment manipulation has had at best a minor effect on overtreatment and only along the distance, not along the time dimension.

[13] Note that the experimenters always paid exactly the price the taxi driver requested. They never gave any tip and taxi drivers never asked for any.

[14] All the p-values reported in the paper refer to two-sided tests.

Table 2

Overtreatment

	CTR	MOH	Overall average
Panel (a): duration index			
Male passengers	1.124	1.135	1.130
Female passengers	1.152	1.126	1.139
Overall average	1.138	1.130	1.134
Panel (b): distance index			
Male passengers	1.056	1.071	1.064
Female passengers	1.053	1.084	1.068
Overall average	1.055	1.077	1.066

Notes. Panel (*a*): The duration index is the ratio of time driven in each ride to time driven in the quickest ride in that particular quadruple. Panel (*b*): The distance index is the ratio of distance driven in each ride to distance driven in the shortest ride in that particular quadruple. CTR refers to the control treatment and MOH refers to the moral hazard treatment.

2.2. *Overcharging and Total Fare*

The values in Table 3 report the frequencies with which drivers overcharged their customers. In total, some form of overcharging was observed in 112 of 400 rides, or in 28% of all cases.[15] In the large majority of these cases (86 out of 112, or 76.8%) bogus surcharges were applied, namely higher-than-justified extras from and to the airport, the port, the railway station and the bus station. The second most frequent source of overcharging (14 cases) were manipulated taximeters or the use of the night tariff during daytime, while rounding-up (not tipping!) of the price (by more than 5%) accounted for the remaining twelve cases of overcharging.

The comparison between the two treatments reveals that overcharging was much more frequent in the moral hazard than in the control treatment (36.5% in MOH *versus* 19.5% in CTR; $p < 0.001$, Fisher's exact test). We view this as a very important finding of our study, because it points towards a statistically significant and economically important causal effect of second-degree moral hazard on market outcomes in the credence goods setting under consideration.

Disaggregating by gender gives a more nuanced impression of this treatment effect. As it turns out, in the control treatment women face overcharging more frequently than men (26% *versus* 13%; $p = 0.031$, Fisher's exact test) but at the same time the charging behaviour towards them is less responsive to our moral hazard manipulation: while the rate of overcharging for women increases by ten percentage points from 26% in treatment CTR to 36% in MOH, the difference is not statistically significant ($p = 0.17$, Fisher's exact test). For male passengers, the difference is even more pronounced – at 24 percentage points – and highly significant (37% *versus* 13%; $p < 0.001$). A consequence of the higher starting value and the weaker treatment effect

[15] In Balafoutas *et al.* (2013), the overall overcharging rate is 11.2% and thus considerably lower. However, since Balafoutas *et al.* (2013) had only male passengers, and those in three different information conditions, the only clean comparison (with identical treatment conditions) is between male non-local natives in Balafoutas *et al.* (2013) – with 7.8% overcharging rate – and male experimenters in the CTR-condition here – with 13% overcharging rate. This difference is not significant according to a χ^2-test ($p > 0.2$).

Table 3

Overcharging and Price

	CTR	MOH	Overall average
Panel (a): overcharging frequency (mean overcharging amount in parentheses, in €)			
Male passengers	0.13 (0.72)	0.37 (1.46)	0.25 (1.09)
Female passengers	0.26 (1.10)	0.36 (1.40)	0.31 (1.25)
Overall average	0.20 (0.91)	0.37 (1.43)	0.28 (1.17)
Panel (b): price index			
Male passengers	1.075	1.153	1.114
Female passengers	1.109	1.177	1.143
Overall average	1.092	1.165	1.129

Notes. Panel (*a*): Overcharging frequency refers to the share of rides that have been classified as cases of overcharging. In parentheses, we report the mean unconditional overcharging amount (which is zero if overcharging has not taken place). Panel (*b*): The price index is the ratio of total price paid in each ride to the lowest total price paid in that particular quadruple. CTR refers to the control treatment and MOH refers to the moral hazard treatment.

for females than for males is that in the presence of moral hazard the rates of overcharging are very similar across gender.

In Table 4, we present the results from a series of regressions on overcharging and price.[16] In all six specifications, standard errors are clustered by quadruple, in order to account for the interdependencies among the four simultaneous rides. We also include route fixed effects, since some routes (in particular the relatively longer ones and the ones from and to the airport) are more susceptible to fraudulent behaviour by drivers.[17] Columns 1 and 2 report probit regressions with the overcharging indicator as the dependent variable. These regressions confirm the above insights with respect to the overcharging frequency, showing a strongly significant effect of treatment MOH on the likelihood of overcharging. Regarding gender, we see in column 2 that women are, *ceteris paribus*, 18.1% more likely than men to face overcharging in the control treatment, while the negative and significant interaction of female passengers and moral hazard is in agreement with our previous observation that the effect of the moral hazard manipulation on overcharging is more pronounced for male than for female passengers.

In addition to overcharging being much more frequent in treatment MOH, the mean amount by which drivers artificially increased the payable fare is also higher in this treatment than in the control (€1.43 *versus* €0.91). The mean overcharging amount by treatment and gender is reported in parentheses in panel (*a*) of Table 3, while in columns 3 and 4 of Table 4 we report Tobit regressions with this amount as the dependent variable. The results show that the moral hazard manipulation leads to significantly higher overcharging amounts compared to the control, and also that – similarly to the overcharging frequency – women face more fraud than men in treatment CTR but not in MOH.

[16] The data file as well as the commands that we have used to run these regressions are available on the ECONOMIC JOURNAL website.

[17] All of the results that are presented in this Section are robust to alternative specifications, such as experimenter fixed effects, time fixed effects or the inclusion of further explanatory variables (driver's age, vehicle type, weather conditions etc.). These additional explanatory variables are always insignificant and lead to a worsening of the Akaike and Bayesian information criteria.

Table 4

Regressions on Overcharging and Price

	(1)	(2)	(3)	(4)	(5)	(6)
	Probit	Probit	Tobit	Tobit	Tobit	Tobit
Dependent variable	Overcharging indicator	Overcharging indicator	Overcharging amount	Overcharging amount	Price index	Price index
Moral hazard	0.193***	0.291***	2.735***	4.033***	0.109***	0.132***
	(0.048)	(0.061)	(0.704)	(1.007)	(0.025)	(0.030)
Female	0.072	0.181***	0.965	2.384**	0.037*	0.062**
	(0.044)	(0.058)	(0.681)	(0.996)	(0.023)	(0.031)
Female × moral hazard		−0.166**		−2.383*		−0.047
		(0.065)		(1.307)		(0.044)
Fixed effects	Route	Route	Route	Route	Route	Route
N	400	400	400	400	400	400

Notes. All reported coefficients are marginal effects. Standard errors in parentheses, clustered by quadruple. The overcharging indicator in (1) and (2) is 1 if overcharging took place, and 0 otherwise. The unconditional overcharging amount in (3) and (4) is left-censored at 0. The price index in (5) and (6) is left-censored at 1. *, ** and *** denotes significance at the 10%, 5% and 1% levels respectively. All p-values and corresponding significance levels are obtained using the Bonferroni correction for multiple hypotheses testing within each regression (but not across regressions).

Given the substantially higher overcharging frequencies and overcharging amounts in the moral hazard treatment compared to the control, it is natural to expect that second-degree moral hazard will have a positive overall impact on the total price paid by passengers, thereby reducing consumer surplus. This is indeed what we observe in the data: panel (*b*) of Table 3 reports – by treatment and gender – the price index, defined in a way analogous to the overtreatment indices (the price paid divided by the minimum price in that particular quadruple). The mean value of the index is substantially higher in MOH than in CTR and the difference is strongly significant ($p < 0.001$; Wilcoxon signed-rank test). Graphically this impression is confirmed in Figure 1, which shows that the cumulative distribution function of the price index in treatment MOH first-order stochastically dominates that in treatment CTR. This information is another key finding of our study, since it implies that consumer expenditures are indeed higher in the presence of moral hazard, *ceteris paribus*. Econometrically, the effect of the moral hazard manipulation on price is analysed in columns 5 and 6 of Table 4, which use Tobit models with the price index as the dependent variable (left-censored at 1). The treatment variable is highly significant, thus confirming our main result. In addition, women pay slightly higher total prices than men, primarily due to the fact that they are more likely to face overcharging in the control treatment.[18]

[18] In addition to the price index, we have also run OLS regressions in which we quantify consumer expenditure with total price paid as our dependent variable. These regressions confirm the highly significant effect of the moral hazard manipulation. A drawback of the variable total price is that it does not take into account the fact that we have routes of different length. This is the reason why we decided to report the results for the price index in the main text.

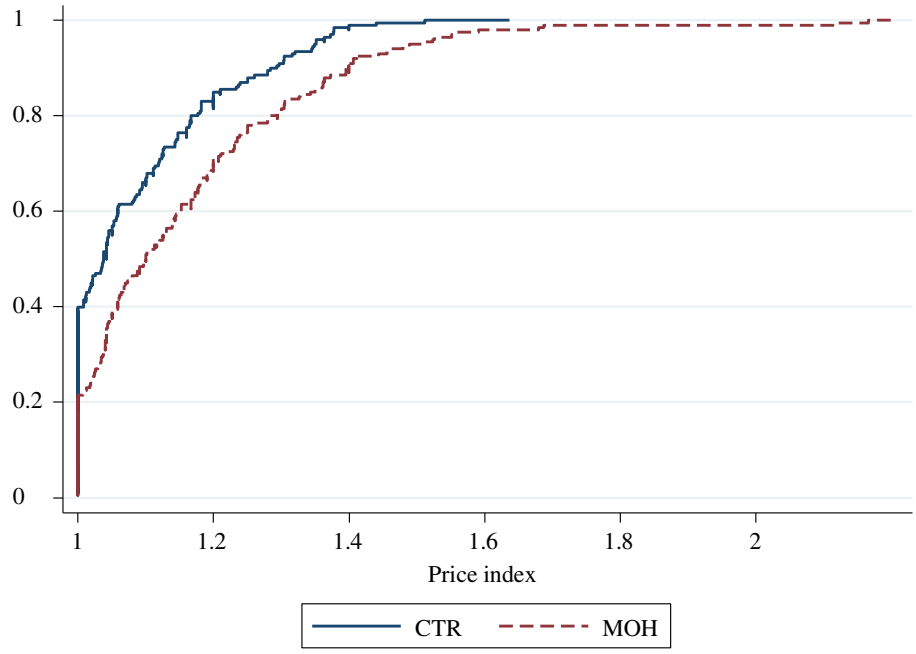

Fig. 1. *Cumulative Distribution Function of Price Index, by Treatment*

3. Discussion

3.1. *Overtreatment Versus Overcharging*

We have documented the presence of strong treatment differences with respect to overcharging behaviour, both in terms of the frequency with which the phenomenon occurs and the magnitude of overcharging. On the other hand, the values of the overtreatment indices are very similar across treatments. Why is it that drivers only change their behaviour in the former fraud dimension but practically not in the latter? One plausible and straightforward explanation is that overcharging is more lucrative and increases the driver's income without any additional costs such as fuel costs, depreciation etc. Therefore, provided a driver believes that his passenger will not complain about a price increase due to fraud in treatment MOH, he should resort to overcharging in order to achieve this price increase.

A further explanation for the observed data pattern relies on the opportunity costs of time. In particular, it is reasonable for a driver to expect that a passenger in the moral hazard treatment – but not in the control treatment – does not mind paying a higher price but that he does mind being taken on a detour because he bears a cost in the form of the longer duration of the ride. Then, a driver who wants to increase income and at the same time reduce the likelihood of being confronted by a passenger should overcharge the passenger but not overtreat him. This line of reasoning is fully consistent with the insignificant difference in the duration index across treatments.

3.2. *Alternative Explanations for the Observed Treatment Differences*

Our main research hypothesis and the experimental design are based on the premise that service provision in credence goods markets is affected by second-degree moral hazard as we have defined it in the introduction. An important assumption behind this hypothesis is that passengers in the MOH treatment are perceived as less likely to exert control over the taxi driver, because they have conveyed to the driver the information that their expenses are reimbursed by their employer. From the perspective of the driver this means that fraudulent behaviour is more likely to go undetected or unreported and, as a consequence, drivers exercise fraud more frequently on average.

A variant of this story could also explain the gender differences in overcharging patterns observed in the data: if female passengers are perceived by – at least some of – the drivers as less likely to complain than male passengers, then those drivers have a more pronounced tendency to overcharge females than males even in the control treatment. The same story might also explain why the treatment difference of adding the phrase 'reimbursed by the employer' is less pronounced for females: if some drivers perceive women as unlikely to complain in general, then adding the phrase does not change their perception. Thus, all the patterns we observe in the data are consistent with such a story. As an anonymous referee has pointed out, one could also model the story by assuming there are three types of drivers:

(*i*) drivers who always attempt to overcharge;
(*ii*) drivers who only overcharge if they think that the passenger will not complain; and
(*iii*) drivers who never overcharge.

By assuming in addition that among the drivers in class (*ii*) a fraction perceives females as less likely to complain in general while the rest does not distinguish between male and female passengers, we can get exactly the pattern we observe in the data. Of course, while our hypothesis seems plausible, we must acknowledge the existence of other stories, which could potentially explain parts of the patterns observed in the data.

3.2.1. *Firms versus individuals: distributional preferences*

It could be argued that drivers are more willing to overcharge passengers who are having their expenses reimbursed by a firm not because of the role of second-degree moral hazard but due to the fact that firms are considered as wealthier than individuals. If this is the case and if drivers are motivated by convex distributional preferences then they might be inclined to charge more when a firm is liable for the payment. This argument could be substantiated by referring to the experimental evidence indicating that distributional preferences are behaviourally relevant in many important market and non-market transactions (see Cooper and Kagel (2012) for a recent survey), and that subjects in the laboratory indeed decide in conformity with convex distributional preferences (Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000; or Charness and Rabin, 2002). While this is a plausible argument, we note, first, that Balafoutas *et al.* (2013) have explicitly tested for the impact of distribu-

tional preferences on service provision in this particular (taxi) market and have failed to find evidence to support such an impact on overtreatment or overcharging decisions. Second, in this experiment we have found a sizeable difference in overcharging between male and female passengers in the control treatment along with a stronger treatment effect for males than for females – a pattern which is hard to reconcile with distributional preferences driving the results.

3.2.2. *Firms versus individuals: morality of fraud and social distance*
Another reason why drivers might be more inclined to overcharge clients who are being reimbursed by their employer is that they feel more comfortable when the victim of fraud is eventually an anonymous entity, as opposed to a specific individual. This could be based either on moral considerations that make a driver more reluctant to steal from an individual than from a legal entity, or on social distance that makes the driver more reluctant to steal from someone who he (or she) directly interacts with than from an anonymous third party.[19] While we cannot entirely rule out such explanations, there are two observations that speak against them. First, they cannot readily account for the gender differences in overcharging patterns and, second, the differences we observe are mainly in the overcharging dimension of fraud while there are hardly any differences in overtreatment across gender and across treatments. Both of these observations are hard to bring in line with a plausible morality (or social distance) story. On the other hand, we have argued that an explanation based on the perceived likelihood that fraudulent behaviour goes undetected or unreported is consistent with the specific pattern we see in the data. In any case, even if those alternative explanations for the treatment differences matter, the implications of our findings would remain the same: market participants who are perceived as having their expenses covered by an anonymous legal person will be more prone to overcharging than those who are paying for the service themselves, resulting in more frequent fraud and higher overall expenditures in the market.

4. Conclusion

We present a field experiment on fraud in a market for a credence good – taxi rides in an unknown city. To the best of our knowledge, ours is the first controlled field experiment to study moral hazard and how it affects market outcomes. Our setting allows us to examine the impact of what we call second-degree moral hazard on the provision of this good. By doing so we study the supply side, rather than the demand side, and how it reacts to moral hazard on the demand side of the market. We consider the latter an important extension of the existing literature because the welfare consequences of this indirect effect of moral hazard might even exceed those of the direct effect. In the context of credence goods markets, this mechanism is particularly relevant because the informational asymmetries present in such markets facilitate supplier-induced demand, or even fraud.

[19] On the effects of social distance on behaviour, see, for instance, Charness *et al.* (2007) or Goette *et al.* (2012).

In our experiment, we find that our moral hazard manipulation has an economically important and statistically significant positive effect on the likelihood and the amount of overcharging, with passengers in that treatment being about 17% more likely to pay higher-than-justified prices for a given ride and payments being about 7% higher on average. This indicates that second-degree moral hazard may have a substantial impact on service provision in a credence goods market. Interestingly, but in line with our hypothesis, we find little or no treatment effect on overtreatment, independently of whether we measure it in terms of duration of the rides or the length of detours taken by taxi drivers.

Our findings are consistent with anecdotal evidence that perceived moral hazard leads to higher bills and thus higher costs in the health care system. For instance, the health care system in Germany is allegedly prone to physicians inflating patients' bills through adding services on the bill that have not been provided. Given that patients do not need to worry about the bill when they are insured because insurance companies reimburse the physicians' bills, this is perhaps not surprising. As a result of this, there is an incentive for the funders (like employers or insurance companies) to reduce the extent of (second-degree) moral hazard by vertical integration with the service provider. For instance, employers could buy their own fleet of cars and run a chauffeur service (with drivers paid with fixed wages) for their employees. In the medical industry, managed care basically consists of such vertical integration: insurers (health maintenance organisations) hire their own doctors and pay them in a way as to contain costs (Ellis and McGuire, 1993).[20] This suggests that the second-degree moral hazard phenomenon investigated in this study is a powerful problem on credence goods markets, which – at least in some sectors – has led to the design of fraud-restraining institutions.

Appendix A. Additional Information on Methods

Table A1
Origins and Destinations

Name	Description
Airport	E. Venizelos International Airport
Glyfada	High-income suburb, southern Athens
Karaiskaki Square	Run-down neighbourhood (central)
Kifissia	High-income residential suburb, northern Athens
Port (Piraeus)	Main commercial and tourist port
Syntagma	Central square, foreigner area
Railway station	Main railway station, all intercity trains (origin only)
Evangelismos	Central Athens
Bus station	Main bus station, services mainly to southern and central Greece
Pagrati	Central residential area (origin only)

[20] We thank an anonymous referee for suggesting this discussion.

Table A2

Routes

#	Origin	Destination
1	Pagrati	Port
2	Port	Karaiskaki Square
3	Karaiskaki square	Kifissia
4	Kifissia	Syntagma
5	Railway station	Port
6	Port	Bus station
7	Bus station	Port
8	Evangelismos	Glyfada
9	Glyfada	Evangelismos
10	Airport	Glyfada
11	Glyfada	Airport

Fig. A1. *GPS Satellite Logger Used During Taxi Rides*

University of Innsbruck
University of Cologne and University of Innsbruck

Accepted: 12 January 2015

Additional Supporting Information may be found in the online version of this article:

Data S1.

References

Akerlof, G. (1970). 'The market for 'lemons': quality uncertainty and the market mechanism', *Quarterly Journal of Economics*, vol. 84(3), pp. 488–500.
Balafoutas, L., Beck, A., Kerschbamer, R. and Sutter, M. (2013). 'What drives taxi drivers? A field experiment on fraud in a market for credence goods', *Review of Economic Studies*, vol. 80(3), pp. 876–91.
Bartling, B. and von Siemens, F. (2010). 'The intensity of incentives in firms and markets: moral hazard with envious agents', *Labor Economics*, vol. 17(3), pp. 598–607.

Bolton, G.E. and Ockenfels, A. (2000). 'ERC: a theory of equity, reciprocity, and competition', *American Economic Review*, vol. 90(1), pp. 166–93.
Camerer, C., Babcock, L., Loewenstein, G. and Thaler, R. (1997). 'Labor supply of New York City cabdrivers: one day at a time', *Quarterly Journal of Economics*, vol. 112(2), pp. 407–41.
Castillo, M., Petrie, R., Torero, M. and Vesterlund, L. (2013). 'Gender differences in bargaining outcomes: a field experiment on discrimination', *Journal of Public Economics*, vol. 99, pp. 35–48.
Charness, G., Haruvy, E. and Sonsino, D. (2007). 'Social distance and reciprocity: the internet vs. the laboratory', *Journal of Economic Behavior and Organization*, vol. 63(1), pp. 88–103.
Charness, G. and Rabin, M. (2002). 'Understanding social preferences with simple tests', *Quarterly Journal of Economics*, vol. 117(3), pp. 817–69.
Chiappori, P.-A., Durand, F. and Geoffard, P.-Y. (1998). 'Moral hazard and the demand for physician services: first lessons from a French natural experiment', *European Economic Review*, vol. 42(3-5), pp. 499–511.
Chiappori, P.-A. and Salanié, B. (2000). 'Testing for asymmetric information in insurance markets', *Journal of Political Economy*, vol. 108(1), pp. 56–78.
Cooper, D.J. and Kagel, J.H. (2012). 'Other regarding preferences: a selective survey of experimental results', in (J.H. Kagel and A. Roth, eds.), *Handbook of Experimental Economics, vol. 2*. Princeton, NJ: Princeton University Press. Available at: http://www.econ.ohio-state.edu/kagel/Other_Regarding%20Preferences_survey.pdf (last accessed 28 March 2015).
Crawford, V.P. and Meng, J. (2011). 'New York City cabdrivers' labor supply revisited: reference–dependent preferences with rational–expectations targets for hours and income', *American Economic Review*, vol. 101(5), pp. 1912–32.
Darby, M.R. and Karni, E. (1973). 'Free competition and the optimal amount of fraud', *Journal of Law and Economics*, vol. 16(1), pp. 67–88.
Dionne, G., Michaud, P.-C. and Dahchour, M. (2013). 'Separating moral hazard from adverse selection and learning in automobile insurance: longitudinal evidence from France', *Journal of the European Economic Association*, vol. 11(4), pp. 897–917.
Doran, E., Robertson, J. and Henry, D. (2005). 'Moral hazard and prescription medicine use in Australia – the patient perspective', *Social Science & Medicine*, vol. 60(7), pp. 1437–43.
Dulleck, U. and Kerschbamer, R. (2006). 'On doctors, mechanics, and computer specialists: the economics of credence goods', *Journal of Economic Literature*, vol. 44(1), pp. 5–42.
Dulleck, U., Kerschbamer, R. and Sutter, M. (2011). 'The economics of credence goods: on the role of liability, verifiability, reputation and competition', *American Economic Review*, vol. 101(2), pp. 526–55.
Einav, L., Finkelstein, A., Ryan, S., Schrimpf, P. and Cullen, M.R. (2013). 'Selection on moral hazard in health insurance', *American Economic Review*, vol. 103(1), pp. 178–219.
Ellis, R.P. and McGuire, T.G. (1993). 'Supply–side and demand–side cost sharing in health care', *Journal of Economic Perspectives*, vol. 7(4), pp. 135–51.
Farber, H. (2005). 'Is tomorrow another day? The labor supply of New York City cabdrivers', *Journal of Political Economy*, vol. 113(1), pp. 46–82.
Farber, H. (2008). 'Reference–dependent preferences and labor supply: the case of New York City taxi drivers', *American Economic Review*, vol. 98(3), pp. 1069–82.
Fehr, E. and Schmidt, K.M. (1999). 'A theory of fairness, competition, and cooperation', *Quarterly Journal of Economics*, vol. 114(3), pp. 817–68.
Goette, L., Huffman, D. and Meier, S. (2012). 'The impact of social ties on group interactions: evidence from minimal groups and randomly assigned real groups', *American Economic Journal: Microeconomics*, vol. 4(1), pp. 101–15.
Grossman, S. and Hart, O. (1983). 'An analysis of the principal–agent problem', *Econometrica*, vol. 51(1), pp. 7–45.
Holmström, B. (1979). 'Moral hazard and observability', *Bell Journal of Economics*, vol. 10(1), pp. 74–91.
Huck, S., Lünser, G. and Tyran, J.–R. (2007). 'Pricing and trust,' Working Paper, University College London.
Huck, S., Lünser, G. and Tyran, J.-R. (2012). 'Competition fosters trust', *Games and Economic Behavior*, vol. 76 (1), pp. 195–209.
Jackson, C.K. and Schneider, H. (2011). 'Do social connections reduce moral hazard? Evidence from the New York City taxi industry', *American Economic Journal: Applied Economics*, vol. 3(3), pp. 244–67.
List, J.A. (2006). 'Field experiments. A bridge between lab and naturally occurring data.' *Advances in Economic Analysis and Policy* 6: Article 8.
List, J.A. and Reiley, D.R. (2008). 'Field experiments', in (S.N. Durlauf and L.E. Blume, eds.), *The New Palgrave Dictionary of Economics*, 2nd edn. Basingstoke: Palgrave Macmillan.
Mookherjee, D. (1984). 'Optimal incentive schemes with many agents', *Review of Economic Studies*, vol. 51(3), pp. 433–46.
Pauly, M.V. (1974). 'Overinsurance and public provision of insurance: the role of moral hazard and adverse selection', *Quarterly Journal of Economics*, vol. 88(1), pp. 44–62.
Rothschild, M. and Stiglitz, J.E. (1976). 'Equilibrium in competitive insurance markets: an essay on the economics of imperfect information', *Quarterly Journal of Economics*, vol. 90(4), pp. 629–49.

Schaller, B. (2005). 'A regression model of the number of taxicabs in US cities', *Journal of Public Transportation*, vol. 8(5), pp. 63–78.
Schaller consulting. (2006). *The New York City Taxicab Fact Book*. Available at: http://www.schallerconsult.-com/taxi/taxifb.pdf (last accessed 1 July 2013).
Schneider, H. (2010). 'Moral hazard in leasing contracts: evidence from the New York City taxi industry', *Journal of Law and Economics*, vol. 53(4), pp. 783–805.
Sülzle, K. and Wambach, A. (2005). 'Insurance in a market for credence goods', *Journal of Risk and Insurance*, vol. 72(1), pp. 159–76.