



Seller Opportunism in Credence Good Markets – The Role of Market Conditions

Katharina Momsen, Markus Ohndorf

Working Papers in Economics and Statistics

2022-10



University of Innsbruck
Working Papers in Economics and Statistics

The series is jointly edited and published by

- Department of Banking and Finance
- Department of Economics
- Department of Public Finance
- Department of Statistics

Contact address of the editor:
Faculty of Economics and Statistics
University of Innsbruck
Universitaetsstrasse 15
A-6020 Innsbruck
Austria
Tel: + 43 512 507 96136
E-mail: Dean-EconStat@uibk.ac.at

The most recent version of all working papers can be downloaded at
<https://www.uibk.ac.at/eeecon/wopec/>

For a list of recent papers see the backpages of this paper.

Seller Opportunism in Credence Good Markets – The Role of Market Conditions*

Katharina Momsen[†] Markus Ohndorf[‡]

June 13, 2022

Abstract

We report the results of an experiment to systematically investigate the influence of different settings in credence good markets on opportunism in the sellers' decisions. We find that, as predicted by a cognitive dissonance model, the specific choice of the design features might be less innocuous than generally presumed: sellers' decisions made under a direct sales regime are significantly more opportunistic than purchase recommendations. Furthermore, average choices are more opportunistic when a costless diagnosis is required to assess the buyer's needs — sellers exploit moral wiggle room by avoiding information. Yet, this effect is only present for purchase recommendations, not direct sales. Both of these effects significantly affect market efficiency. Generally, the parametrization of the decision problem has a strong influence on opportunism, as predicted. Here, we find that sellers tend to overtreat and buyers self-select into overtreatment.

JEL Classification: C90, D47, D82, D91, L15

Keywords: information avoidance, credence goods, moral wiggle room, norm activation model, online experiment

*This study was preregistered at AsPredicted (https://aspredicted.org/LDC_28G) and was approved by the Board for Ethical Questions in Science of the University of Innsbruck (certificate 37/2021).

[†]University of Innsbruck, Innsbruck, Austria, katharina.momsen@uibk.ac.at

[‡]University of Innsbruck, Innsbruck, Austria, markus.ohndorf@uibk.ac.at (corresponding author)

1 Introduction

Over the last decade, economists have intensified their endeavor to provide experimental evidence on whether markets erode moral behavior. Yet, the experimental evidence for moral behavior in market contexts is mixed. While [Falk and Szech \(2013\)](#) and, to a certain extent, [Bartling et al. \(2015\)](#) find that markets may undermine moral values, other studies find only limited support for this claim ([Kirchler et al., 2016](#); [Pigors and Rokenbach, 2016](#); [Sutter et al., 2020](#)). One argument for markets undermining morals is based on the perception that in market transactions the other side remains more or less anonymous, which is certainly the case in large marketplaces, like stock markets ([Sandel, 2012](#)). However, many types of transactions require direct interaction between the supply and the demand side of a market. This holds, in particular, for markets for credence goods, which often involve the provision of goods or services that are consumer-specific. In these contexts, one might expect that moral preferences influence the sellers' behavior, as the counterpart is more likely to be perceived as a human-being.

Indeed, as experimentally shown in [Kerschbamer et al. \(2017\)](#), sellers in such markets often deviate from the neoclassical assumption of perfectly rational, purely self-interested decision-makers and exhibit a behavior that seems to be influenced by their social preferences.¹ Yet, beyond these insights, it remains unclear under which circumstances sellers tend to refrain from opportunism in markets for credence goods. In this paper, we investigate if different market conditions discussed in the experimental literature on credence goods affect the sellers' moral behavior, even if anonymity is guaranteed.

Generally, a multitude of products and services fall into the category of credence goods ([Darby and Karni, 1973](#)). Prominent examples are visits to the doctor or the car mechanic, taxicab rides in unknown cities, complex products and financial advice. The common feature of these products and services is the large information asymmetry between supply and demand side. The typical setting for a credence good is as follows: a consumer intends to purchase a solution to cover a specific need. The solution can be large or small, with the consumer not knowing which is adequate to fulfill her needs. The seller, on the contrary, can observe the buyer's needs and knows which kind of treatment or product is most adequate to address them. Similar to the case of experience goods ([Nelson, 1970](#)), *ex ante*, the consumer does not know what kind of service, treatment or product she needs. However, in contrast to experience goods, the consumer of a credence good cannot properly assess even *ex post*, i.e. after the purchase has taken place, if she has received the appropriate treatment or purchased

¹In markets for label credence goods where choices also affect third parties (see, [Balafoutas and Kerschbamer, 2020](#), for a definition of the term), one would expect the morality of decisions to be driven by the demand side ([Momsen and Ohndorf, 2020b, 2022](#)).

the correct version of the product. This gives rise to three different types of opportunistic (or even fraudulent) behavior in markets for credence goods (see, e.g., [Dulleck and Kerschbamer, 2006](#); [Dulleck et al., 2011](#)): overtreatment, undertreatment and overcharging.² All of these types of mistreatment represent an exploitation of asymmetric information and hence yield inefficiencies in the market from which the consumer suffers.

Note that within the literature on credence goods, four different setups coexist that are usually not explicitly delimited. These setups vary with respect to how the seller's choice is implemented for the buyer and whether a diagnosis is necessary. For the former variation, (i) the seller might either give a recommendation and leave it to the consumer to decide if she chooses the recommended option or not, or (ii) he might directly implement a solution within a wider contractual agreement. The second variation concerns the seller's diagnosis: (a) either it is directly obvious to the seller what type of product would fit the buyer best, or (b) the seller needs to acquire information to understand the buyer's situation. Assuming perfectly rational decision-makers with stable preferences on both market sides, the four scenarios should yield identical market outcomes. If we depart from this assumption, however, it is unclear how the decision frame influences the seller's choices. In this paper, we propose an experiment to systematically investigate the influence of the decision setting on outcomes in credence goods markets.

Let us first consider the variation with respect to the mode of acquisition of the credence good. An example for variety (ii), a situation where the seller directly implements a solution, is a standard taxicab ride. Instead of recommending a specific route, a cab driver will choose their preferred route and take the customer to their destination. Similarly, a car mechanic asked to "get the car back on the road" does not give a recommendation on what should be repaired, but instead directly replaces the broken parts. The same applies to craftsmen such as plumbers asked to "just fix" a lavatory or an IT consultant to "just get the system running". Variety (i), where the seller gives a recommendation for the buyer to accept, is also very common. It certainly arises in purchases of complex products, where consumers often rely on advice by sales personnel.

In a model with rational types (altruists and opportunists), sellers' choices will be identical for both varieties. However, this equivalence is challenged if we consider the behavioral effect of additional decision features on moral decisions. For example, giving an opportunistic recommendation that does not serve the buyer's needs might be considered a lie. In this case, an aversion to lie, as identified in [Gneezy \(2005\)](#), would

²While in the case of overtreatment, the consumer's utility would have been largest when receiving the small treatment, she has been provided the larger, more expensive version. In the case of undertreatment, she has received a treatment that does not solve her problem. Finally, in the case of overcharging she has received the small version of the treatment, but paid for the large version.

lead to a lower amount of opportunism for recommendations (variety (i)) than for direct sales (variety (ii)), as the latter does not involve an active act of lying.³ Perceiving a buyer as an active decision-maker may also increase the salience of the impact of his choice on another individual and thus trigger a more altruistic (hence less opportunistic) decision (Bandura, 2002; Petras et al., 2016). This would be consistent with models in the tradition of Schwartz (1977), according to which social norms need to be activated within the individual to result in moral or pro-social behavior.

A similar question arises with respect to the sellers' investment in a diagnosis of the buyer's needs. For example, a doctor typically performs a diagnosis before he gives a recommendation for a treatment, which corresponds to variety (b) above. Yet, it may also be the case that symptoms are so specific that he, as an expert, immediately knows which kind of treatment will cure his patient (variety (a)). Similarly, a car mechanic usually needs to examine a car closely to understand potential issues (variety (b)). Yet, there are certainly also situations where, for him as a specialist, the appropriate repair is directly obvious (variety (a)).

Again, within standard microeconomics if the diagnosis is without cost, both types of situations are considered equivalent and should lead to the same decision on the part of the seller (Bolton and Dewatripont, 2005). Yet, there is a growing literature that indicates that in moral contexts, subjects tend to avoid information in order to exploit moral wiggle room (Dana et al., 2007; Grossman and van der Weele, 2017). If this phenomenon is relevant in our context, then one would expect that sellers may act more selfishly in situations of variant (b), than when they can immediately observe which treatment would suit the buyer best (i.e. variant (a)). These considerations are distinct from analyses of diagnosis in information economics, as the costs of information are zero here.

In this paper, we report the results of an online experiment where subjects interact in the roles of buyers and sellers to study the effects of the four variants of credence good markets on seller opportunism described above in four different treatments. Moreover, we vary the parametrization determining if buyers' and sellers' interests are aligned or conflicting, as well as the relative payoff differences of buyers and sellers (weak vs. strong conflict of interests vs. symmetric situation). The focus is, hence, on investigating the sellers' tendency to choose opportunistically via under- or overtreatment of the buyer.

To derive our hypotheses, we extend the cognitive dissonance model developed in Momsen and Ohndorf (2022) which is, in turn, based on Rabin (1994). The model is to reflect the 'internal' information economics of the seller. This allows us to better

³In the literature on lying aversion, this is often attributed to (and represented as) additively separable 'costs of lying', solely dependent on differences in own, or both, payoffs (e.g. Kartik, 2009; Gneezy et al., 2013; Dellavigna et al., 2017; Kajackaite and Gneezy, 2017; Khalmetski and Sliwka, 2019).

identify the conditions for which self-serving information avoidance occurs, and is suitable to integrate a potential activation of personal altruistic norms. As it turns out, such a cognitive dissonance model seems to be better suited to predict the experimental results than considerations based on additively separable costs of lying.

We find that opportunism is not ubiquitous, but occurs on average in about 50% to 70% of sellers decisions, depending on the treatment. The share of opportunistic decisions is indeed significantly higher with direct provision of products than with purchase recommendations. Furthermore, average recommendations under hidden but revealable information are also more opportunistic than those made under full information. Hence, sellers exploit moral wiggle room. Yet, interestingly, we do not find evidence of such self-serving information avoidance when their decisions are directly implemented for buyers. This might be explained by the even greater normative distance to a completely passive anonymous buyer, such that norm activation is lower for direct sales than with recommendations. Giving a recommendation obviously involves a human buyer that ultimately takes a purchase decision, which might lead to a higher tendency to activate altruistic norms.

We also observe that both buyers and sellers react to the parametrization: sellers' choices are more selfish in case of a weak conflict of interests and buyers have a tendency to choose the option with the lower variance in payoffs. Hence, in the terminology of the credence goods literature, sellers tend to overtreat and buyers have a tendency to self-select into overtreatment.

The remainder of this paper is structured as follows: the following section provides a literature overview to set the different variants into the context of the literature. The experimental design is then described in Section 3. Section 4 presents the behavioral predictions based on a simple model and Section 5 presents the results. The last section concludes. Screenshots of the decision screens including the instructions are relegated to the Appendix.

2 Related Literature

The purpose of this literature section is two-fold: firstly, we seek to organize the existing literature on credence goods with respect to the market characteristics altered in the experiment, which correspond to the four varieties discussed in the introduction. Secondly, we discuss the literature in behavioral economics that is related to this research.

Credence goods with recommendation without diagnosis (variety (i) and (a))

In this literature, it is assumed that the buyer has the option to accept or reject the recommended treatment. If she refuses to undergo the recommended treatment, no trade

takes place and she and her seller end up with the outside option. Most contributions to this literature are theoretical in nature, e.g. [Inderst and Ottaviani \(2009, 2012b\)](#); [Liu \(2011\)](#); [Fong et al. \(2014\)](#); [De Jaegher \(2012\)](#); [Pitchik and Schotter \(1987\)](#); [Fong \(2005\)](#); [Sülzle and Wambach \(2005\)](#). Yet, there is also an increased interest from the experimental community in this type of good (see, for example, [Schneider et al., 2021](#)). While most studies in this strand deal with abstract settings, financial advisory services are a common motivating example.

Credence goods with recommendation with diagnosis (variety (i) and (b))

In a different setup common in the literature on credence goods, the seller needs to perform a diagnosis to observe the buyer's needs before he gives a recommendation. The seminal paper by [Darby and Karni \(1973\)](#) falls into this category, as well as [Dulleck and Kerschbamer \(2006, 2009\)](#); [Wolinsky \(1993, 1995\)](#); [Emons \(1997, 2001\)](#); [Alger and Salanié \(2006\)](#). [Anagol et al. \(2017\)](#) run a field experiment in the Indian market for life insurance. In this setting, advisors clearly need to invest some effort into finding out which kind of insurance would benefit their client most and thus be able to give a trustworthy recommendation. In the laboratory experiment presented in [Momsen \(2021\)](#), buyers have the possibility to choose an option that the seller has not recommended, instead of following the advice or choosing the outside option. Arguably, this implementation may bring the experimental setup a little closer to the world outside the lab, as actual consumers who are in need of a certain product might choose a similar product provided by the same seller, which is usually the case for purchases of white goods in a specialized store.

Credence goods without recommendation without diagnosis (variety (ii) and (a))

The literature where the seller directly provides the product, service or treatment to a completely passive consumer is mostly based on experiments. The large-scale study by [Dulleck et al. \(2011\)](#) features one of the earliest of these lab experiments, which they implemented in an abstract setting. Note that the experimental setup in [Dulleck et al. \(2011\)](#) diverges in two ways from the one in the closely related theoretical paper by [Dulleck and Kerschbamer \(2006\)](#): experimental sellers do not give treatment recommendations, but directly provide the treatment. Moreover, they do not need to undergo a diagnosis to observe the buyers' needs. Other studies that rely on (field) experimental methods also fall into the category of direct provision without initial diagnosis, see e.g. [Balafoutas et al. \(2013, 2017\)](#); [Beck et al. \(2014\)](#); [Kerschbamer et al. \(2017\)](#); [Hennig-Schmidt et al. \(2011\)](#); [Hennig-Schmidt and Wiesen \(2014\)](#); [Godager and Wiesen \(2013\)](#); [Mimra et al. \(2016\)](#). While [Balafoutas et al. \(2013, 2017\)](#) perform field experiments investigating fraud in the market for taxicab rides, [Hennig-Schmidt et al. \(2011\)](#); [Hennig-Schmidt and Wiesen \(2014\)](#); [Godager and Wiesen \(2013\)](#) investigate the

market for health services experimentally with medical students representing the supply side.

Credence goods without recommendation with diagnosis (variety (ii) and (b))

There are not many contributions dealing with credence goods with direct provision of the service *and* necessary diagnosis. Prominent examples are the field experiments presented in [Kerschbamer et al. \(2016, 2019\)](#), which investigate computer repairs, as the experts need to perform some, albeit simple diagnosis to figure out what is preventing the computers from booting.

Credence goods and social preferences

Experimental evidence in [Kerschbamer et al. \(2017\)](#) speaks against the notion of sellers' being perfectly self-interested in credence goods markets. This study finds that less than a fourth of the subjects act according to the standard assumption of sellers being rational maximizers of their own payoff. They conclude that these social preferences explain why experimental credence goods markets without verifiability of treatment perform better than predicted by theory. In a similar vein, [Hennig-Schmidt et al. \(2011\)](#) show that monetary considerations cannot be the only motivation for decisions made by subjects on the supply side of their experiment. Using the same experimental data set, [Godager and Wiesen \(2013\)](#) investigate heterogeneity in altruism among subjects. While about a fourth of the subjects values their own payoff more than the payoff of the patient, about 30% attach equal weights to both payoffs and the remaining subjects put a higher weight on the payoff of the patient. [Green \(2014\)](#) finds that retrospective payment schemes crowd out intrinsic motivations. Hence, in the absence of incentive-compatible payment schemes, subjects are more likely to exhibit altruistic behavior.

There is also a sizeable number of theoretical papers that depart from the assumption of opportunistic sellers. [Liu \(2011\)](#), for example, introduces conscientious experts who also derive utility from repairing the consumer's problem, while [Fong et al. \(2014\)](#) introduces honesty in expert decisions and [Inderst and Ottaviani \(2012a\)](#) model advisors to care for the suitability of their recommendation for the customer's needs.

Experts, information, and behavioral effects

When sellers exhibit some sort of social preferences, they may also engage in behavior that allows them to preserve their self-image while maximizing their payoff. For example, [Gneezy et al. \(2020\)](#) show that the experts' advice is sensitive to the timing of information at which potential conflicts of interests are revealed, indicating that such advice can be subject to motivated reasoning ([Epley and Gilovich, 2016](#); [Bénabou and Tirole, 2016](#)). Alternatively, experts may be tempted to simply avoid information on the buyer's needs, thus exploiting moral wiggle room – a phenomenon first described

in the seminal paper by [Dana et al. \(2007\)](#). The rationale behind this motivation for self-serving ignorance is as follows: if a decision-maker with pro-social preferences directly observes the third-party consequences of her choices, she will choose altruistically. If, however, she has the option to remain ignorant about said consequences, she will do so and choose egoistically, as she effectively avoids potential psychological costs from a bad conscience. In other words, she will exploit the wiggle room via information avoidance in order to maximize her monetary payoff while still maintaining a positive self-image.

Many papers (e.g. [Grossman, 2014](#); [Grossman and van der Weele, 2017](#); [Conrads and Irlenbusch, 2013](#); [Kajackaite, 2015](#); [Momsen and Ohndorf, 2020a, 2022](#)) have succeeded in replicating the findings presented in [Dana et al. \(2007\)](#), while others could not find evidence of willful ignorance on the *consumption side* in market experiments ([Bartling et al., 2015](#); [Pigors and Rockenbach, 2016](#)). In contrast, with credence goods, the problem of potential opportunism arises on the *supply side*, as the seller can easily exploit the large information asymmetry in such markets. Thus, any statement on morals influencing market outcome should focus on the supply side, which is the primary motivation of the analysis presented here.

3 Experimental Design

The design of our online experiment is held simple. We examine potential differences in the seller's (he) decision if he directly observes which product fits the buyer's (she) needs or if he actively needs to reveal the respective payoff of the buyer. This variation corresponds to credence good varieties (a) and (b). In addition, we study if the buyer's action space affects the seller's decision, namely whether the fact that the buyer has the option not to follow the seller's advice influences the seller's behavior (corresponding to varieties (i) and (ii)).

In the experiment, the roles of buyers and sellers are randomly assigned and subjects keep these roles throughout the entire experiment. Sellers own two products, A and B, that, if sold, determine the sellers' payoff. The buyers' payoff depends both on the purchased product, A or B, and on the state of the world, H or L. Both states of the world are equally likely and only affect the buyer's payoff, while the seller's payoff for each option is independent of the state of the world. Sellers and buyers make a single transaction in each of the overall ten rounds. In each round, they are assigned a new anonymous counterpart. In other words, there is no competition in the market and buyers cannot choose their respective seller. Furthermore, subjects do not have the option not to interact: in each round, sellers sell exactly one of the two products and buyers purchase one product. While buyers never observe the state of the world – they can never identify which product maximizes their utility and hence should be

purchased – for sellers the observability of the state of the world depends on the treatment. Also, depending on the treatment, the seller either recommends or directly sells one of the two available products. In the former case, the buyer has the possibility to follow the recommendation or to switch to the alternative option. In the case of a direct sale, the buyer simply buys the product selected by the seller.

Products and state of the world vary from round to round, and there is no feedback between the ten rounds. Hence, buyers do not learn if the recommendation they received in the previous round maximized their payoff, nor do sellers learn if the buyer followed their advice. Ultimate payoffs to the subjects are only determined after the completion of all ten rounds. The state of the world also remains hidden to buyers as well as to uninformed sellers in the treatments where the state of the world is hidden, but revealable. The ten rounds contain situations where the product that maximizes the seller’s payoff also maximizes the buyer’s utility (*Aligned Interests*) and situations where the opposite is the case, i.e. the product that maximizes the seller’s profits is associated with the lower level of utility for the buyer (*Conflicting Interests*). Both types of situations are equally likely.

3.1 Treatment Variations

Treatments are implemented between subjects in a 2x2 design, varying the observability of the state of the world as well as how sellers’ decisions are implemented. In the FULLINFO treatments (variety (a)), sellers immediately observe the state of the world. Hence, they know which option maximizes the buyer’s payoff and if there is a conflict of interests between their own and the buyer’s payoff. In the HIDDENINFO treatments (variety (b)), in contrast, sellers initially do not know the state of the world, but they can become aware of the state of the world by clicking a button and thus find out which product maximizes the buyer’s payoff and if there is a conflict of interests. Clicking the button is costless.

The second treatment variation applies to the way how the seller’s decision is implemented for the buyer. In the RECOMMEND treatments (variety (i)), the seller gives a recommendation which the matched buyer can either decide to follow, or choose the alternative product option. Hence, the buyer still has the option to choose, yet without knowing if the recommendation is beneficial to her. To exclude any type of social learning effects, the seller in the RECOMMEND treatments makes his recommendations without knowing which option has been chosen by the matched buyer in the previous rounds. In the SELL treatments where the seller’s choice is sold immediately (variety (ii)), the buyer remains completely passive and the seller’s action directly determines the buyer’s payoff.⁴

⁴See Section B in the Appendix for screenshots of example decision situations in the different treatments. The instructions can be found in Section A.

Table 1: Parametrization

	Seller's payoff	Buyer's payoff in 'L'	Buyer's payoff in 'H'
Strong Conflict of Interests			
Product A	100	100	100
Product B	120	120	20
Product A	90	100	100
Product B	130	120	20
Weak Conflict of Interests			
Product A	100	20	120
Product B	120	100	100
Product A	90	20	120
Product B	130	100	100
Symmetric			
Product A	100	20	100
Product B	120	100	20

Furthermore, we vary the parametrization in a within-subjects manner, implementing situations with a strong conflict of interests, a weak conflict of interests, and symmetric decision situations. In situations with a strong conflict of interests ('StrongCI'), the product maximizing the seller's payoff is associated with a very low payoff for the buyer. For the opposite case, with aligned interests, the buyer benefits about as much as the seller from the seller choosing the optimal product.

In situations with a weak conflict of interests ('WeakCI'), sellers sacrifice only a small amount of the buyer's payoff when maximizing their own payoff. For the opposite situation with aligned interests, buyers benefit disproportionately from the seller taking his own-profit-maximizing option.

Finally, in symmetric decision situations, buyers benefit just as much from an altruistic choice of the seller as they would suffer from an opportunistic choice in a state of the world that represents a situation with conflicting interests (see Table 1 for an overview over example parametrizations).

For the seller's payoff in situations with a 'WeakCI' and a 'StrongCI', we use two different parameter sets, denoted with 'HighPD' for the 130/90-combination and 'LowPD' for 120/100.

For each of these parametrizations, one situation with state of the world 'L', and one with realization 'H' were implemented. Thus, there were 5 situations where interests were aligned, and 5 with conflicting interests, totaling in 10 decisions to be taken within this experiment. The order in which participants were confronted with the 10 decision situations was determined randomly for each participant. Notice that in all

decision situations, except for ‘WeakCI’, the opportunistic choice will be Kaldor-Hicks inefficient if interests are conflicting.

3.2 Experimental Procedure

We used oTree (Chen et al., 2016) to program the experiment, which was implemented with a sample recruited via Prolific.org. We first ran the sessions for the sellers and later on matched their decisions to buyers. One of the ten decision rounds was randomly selected for payment. Following in the steps of previous literature investigating information avoidance in online experiments (see, e.g., Exley and Kessler, 2021), we recruited 100 sellers and 100 buyers per treatment such that, in total, 800 subjects participated in our experiment. Sessions lasted on average 8 minutes and subjects earned an average payoff of \$2.10. As buyers in the SELL treatments remained completely passive, they took part in a choice experiment that is unrelated to this paper.

4 Behavioral Predictions

As we are interested if the different credence good varieties have an influence on opportunism in sellers’ decisions, we restrict our analysis to situations where moral judgement can affect behavior. Hence, throughout the paper, we exclusively analyze situations where interests between buyer and seller are conflicting, as information asymmetries are not exploited in aligned-interest situations. Furthermore, for the following considerations, we assume that spite does not play a significant role in the subjects’ decisions. This is plausible, as anonymous buyers and sellers are re-matched in each round and the results of all transactions are only communicated after the experiment. These features also eliminate any possibility for retaliation, reciprocity, and the build-up of reputation. Thus, by design, any mechanism under which a purely self-interested player would *not* choose opportunistically is excluded.

In this section, we first discuss the predictions of the standard case with perfectly rational agents and altruistic/opportunistic seller types for SELL and RECOMMEND situations. Subsequently, we discuss the notion of lying aversion and norm activation that might alter these predictions. We then present a simple cognitive dissonance model to represent the potential behavioral interactions for all four of our treatments.

4.1 Prediction for complete rationality with altruistic types

We first consider the treatment variations RECOMMEND vs. SELL. Buyers in the SELL treatments are completely passive, whereas buyers in the RECOMMEND treatments can make the active choice to follow the seller’s recommendation, or not. Yet, with credence goods, buyers never observe the state of the world, i.e. even *ex post* they cannot

know which product maximizes their payoff. In this case, the prediction for a rational, completely self-interested seller depends on how much he expects the buyer to trust his recommendation. If sellers anticipate that buyers do not trust the recommendation at all, the recommendation is to be considered an empty signal. In this case, sellers would randomize, in aggregate, over what option they recommend, as any structured choice would not yield additional gain. Thus, if all agents are perfectly rational, and self-interested in the traditional sense, the prediction is complete randomization over both recommendations.

If, however, there also exist altruistic types of sellers, buyers can attribute positive levels of trust towards the seller's recommendation. This would lead to the following behavior: a purely self-interested seller will never choose the option that does not maximize his payoff, as he can expect the buyer to follow his advice with some positive probability. Thus, observing a recommendation that maximizes the seller's payoff, the buyer can neither infer the signal to be truthful, nor the opposite since she is unaware of the state of the world. Hence, the buyer cannot use the signal for any additional inference, and should choose according to her expectations that are exclusively determined by her level of trust. If the opposite signal is given, however, the buyer can indeed make such an inference. Recall that buyers are aware that self-interested sellers never choose the option that does not maximize their payoff. Thus, upon observation of a recommendation that does not maximize the seller's payoff, buyers can directly infer that it is truthful, and made by a seller with pro-social preferences. Buyers will hence always follow such a recommendation.

In this case, we have perfect separation of types when comparing both recommendations exclusively in conflict situations. In this type of situations and for perfectly rational players with stable preferences, we can predict that purely self-interested sellers will always recommend the option that maximizes their payoff. The other option, which is not payoff-maximizing, will be exclusively recommended by sellers of the altruistic type. Thus, with simple binary types and no additional behavioral effects, the share of self-serving choices in situations with conflicting interests in the RECOMMEND treatments should correspond to the one in the SELL treatments.

4.2 Cost of Lying vs. Norm Activation

It is, however, quite likely that additional behavioral effects are at play. Note that if we completely abstract from context, our FULLINFO situations are quite similar to the ones implemented in [Gneezy \(2005\)](#) to investigate the propensity to lie in a Cheap Talk Sender-Receiver Game. They find that the propensity to lie increases in the gains from lying and decreases in the size of the loss to the other party. Thus, in a situation where a lie would harm the other party, self-serving recommendations are less frequent than

self-serving choices in situations where a player directly decides over a specific allocation between both players. In our setup, sellers only have the opportunity to lie in the RECOMMEND treatments, while in the SELL treatments their decision is directly implemented without any additional intervention by the buyer. Thus, if lying aversion plays a similar role as in [Gneezy \(2005\)](#), we would expect a lower share of self-serving recommendations in conflict situations than for sales decisions under the SELL treatments.

In the literature following [Gneezy \(2005\)](#), the aversion to lie is often attributed to (and represented as) additively separable cost of lying (e.g. [Kartik, 2009](#); [Gneezy et al., 2013](#); [Dellavigna et al., 2017](#); [Kajackaite and Gneezy, 2017](#); [Khalmetski and Sliwka, 2019](#)). Such additively separable costs of lying are conceived of as a function of the “size of the lie”, i.e they are dependent on the payoff differences associated with the true and false statements for one or both players. In our experiment, we measure these differences by use of the parametrization WeakCI/StrongCI/Symmetric for payoffs of the buyer and HighPD/LowPD for the seller. These parameters should have different effects on recommendations if costs of lying were to explain treatment differences.

Costs of lying are usually modeled as being independent of the moral norm that co-determines the amount of altruism in a specific situation. Thus, if situational contexts remain the same, as usually investigated in the literature on lying, models using additively separable costs of lying are indeed consistent with the observed behavior. This might, however, not hold if a set of situations is associated with the same payoffs, but different altruistic norms. For example, [Gneezy \(2005\)](#) argues that the propensity to lie to an individual might be lower than to a company. In this case, the norm that is underlying altruistic behavior is different in both cases. Note that this difference in norms is at the basis of the argument of the amorality of anonymous markets in contrast to personal interactions. While participants are anonymous in all treatment variations, our SELL treatments, with their completely passive buyer, might still be perceived as normatively different situations than our RECOMMEND treatments.

As the literature on computer-mediated interaction shows, active interaction within a computerized setup can indeed reduce the effect of dehumanization, which is commonly observed in computer-mediated decision situations ([Cherry et al., 2002](#); [Sassenberg et al., 2005](#); [Bae, 2016](#); [Scholl et al., 2020](#)). In this literature, it is often hypothesized that interaction increases the so-called private self-awareness (e.g. [Sohn et al., 2019](#)), which, in turn, leads to an increase in the feeling of responsibility for one’s own actions (see, e.g., [Lerner and Tetlock, 1999](#)).

To take such potential differences in abiding to personal altruistic norms into account, we borrow the term “norm activation” from the Norm Activation Model by [Schwartz \(1977\)](#) to describe the process leading to different levels of moral engagement ([Bandura et al., 1975](#); [Califano et al., 2022](#)). This approach seems adequate here

as in the Norm Activation Model the three determinants for altruistic behavior co-depend as follows: awareness of the consequences (i.e. processed information) is the antecedent of ascription of responsibility which, in turn, is an antecedent of the activation of personal altruistic norms which are then translated into behavior (De Groot and Steg, 2009). Moral behavior is predicted to occur only if all of these determinants reach their threshold levels (De Groot and Steg, 2009; Onwezen et al., 2013; Blasch and Ohndorf, 2015). Clearly, this is more likely to arise if information on the effects is a) actually revealed, and b) if third-party effects are clearly attributed to agents perceived as human-beings. Given the insights from research on computer-mediated interactions, it seems plausible that the latter is more likely in credence good situations with an active buyer, as in our RECOMMEND treatments. Under SELL, in contrast, the passive buyer setup is closer to an anonymous, impersonalized market situation.

Similar to additive costs of lying, norm activation would hence predict a larger share of altruistic recommendations under RECOMMEND than under SELL. Yet, for the other treatment comparisons, predictions might diverge for both approaches. To assess these differences in a structured manner, we discuss these in the context of a simple cognitive dissonance model extending the considerations in Momsen and Ohndorf (2022), which itself is based on Rabin (1994) and Konow (2000). The model is designed to reflect the ‘internal’ information economics that determine the decision to avoid information, while taking potential differences in norm activation into account.

4.3 Predictions from a simple cognitive dissonance model

4.3.1 Full Information

Consider the situation at the basis of our experiment. Denote with $v \in \{\underline{v}, \bar{v}\}$ the buyer’s payoff in a conflicting interests situation. In such a situation, the product option associated with \underline{v} yields the higher payoff to the seller, denoted with \bar{p} . Analogously, the option associated with the sellers’ lower payoff \underline{p} yields the higher payoff \bar{v} to the buyer. Let j be an index variable, denoting the type of sales situation, with $j = s$ for a SELL situation and $j = r$ for a RECOMMEND situation. To represent an altruistic norm, we define a binary function η_j , with domain $\{\underline{v}, \bar{v}\}$. If this norm is activated in context j , it can potentially influence behavior and takes the value $\eta = \bar{v}$. A norm that is not activated takes the value $\eta = \underline{v}$.

Under full information with conflicting interests, a cognitive dissonance will only occur if individuals with (sufficiently intensive) altruistic preferences choose the self-serving product option $\{\bar{p}, \underline{v}\}$. Following Rabin (1994), we denote the costs from this dissonance when choosing option i with $\Phi(\eta - v_i; \alpha)$, with $\Phi(0) = 0$, $\Phi' > 0$, $\Phi'' > 0$, and parameter α representing the intensity of the altruistic preference. Hence, Φ represents the cost from being inconsistent with own ideals. Under complete information, the

seller's valuations of both product options are in the SELL treatment

$$U_{\neg a} = \bar{p} - \Phi(\eta_s - \underline{v}),$$

$$U_a = \underline{p} - \Phi(\eta_s - \bar{v} = 0) = \underline{p}.$$

Thus, under complete information, the altruistic option a strictly dominates the self-serving option $\neg a$ iff

$$(1) \quad \Phi(\eta_s - \underline{v}) > \Delta_p = \bar{p} - \underline{p}$$

Note that our parametrization (WeakCI/StrongCI/Symmetric) affects the left-hand side of this condition, and the payoff difference to the seller, that can vary in the pair HighPD/LowPD, corresponds to the right-hand side. Thus, the model predicts the largest shares of non-opportunistic behavior for the parametrization StrongCI with LowPD and Symmetric, which is intuitive.

As outlined in Section 4.1, for a RECOMMEND situation, the expected valuation of the seller depends on the buyer's trust in the recommendation. While for completely rational players, we can assume that buyers have no reason to doubt if $\neg a$ is recommended, this does certainly not hold for a recommendation that is advantageous to the seller, i.e. a . Thus, the expected value of the seller's monetary payoff is dependent on how trustworthy he appears to be to the buyer. We denote with $E[p \mid a]$ the seller's expected payoff when recommending the altruistic option a , and with $E[p \mid \neg a]$ the seller's expected payoff when recommending self-serving option $\neg a$. The expected valuations of both options are hence:

$$(2) \quad EU_{\neg a} = E[p \mid \neg a] - \Phi(\eta_r - \underline{v}),$$

$$(3) \quad EU_a = E[p \mid a] - \Phi(\eta_r - \bar{v} = 0) = E[p \mid a].$$

Hence, under complete information, the altruistic option a strictly dominates the self-serving option $\neg a$ iff

$$(4) \quad \Phi(\eta_r - \underline{v}) > \Delta_{E[p]} = E[p \mid \neg a] - E[p \mid a]$$

As laid out above, we can expect $\Delta_{E[p]} \leq \Delta_p$. This needs to be taken into account when comparing the conditions for both treatments that we derived above.⁵ For the sake of simplicity in comparisons, we define $z = \Delta_p - \Delta_{E[p]}$ and reformulate (4) as follows:

⁵We did not elicit the seller's beliefs over his payoffs, as this would likely lead to unwanted priming effects. Such an elicitation would also not add much to our predictions, as the argument made here holds qualitatively for any type of rational expectation.

$$(4') \quad \Phi(\eta_r - \underline{v}) + z > \Delta_p$$

Thus, even if the cost of not adhering to the personal norm $\Phi(\eta - \bar{v})$ remains the same in both situations (i.e. for $\eta_r = \eta_s$), condition (4') is likely to be laxer than condition (1). This leads to the prediction that the share of self-serving decisions under FULLINFO should be larger in the SELL treatment than in the RECOMMEND for any given parametrization combining WeakCI/StrongCI/Symmetric and HighPD/LowPD. The difference in opportunistic behavior is even stronger if we assume $\eta_r > \eta_s$. This would represent the conjecture that some sellers might dehumanize the buyer in the SELL treatment, in which case the norm will not be activated ($\eta_s = \underline{v}$). For such an individual, $\Phi(\eta_s - \underline{v}) = 0$, and condition (1) would never hold. We can hence confidently state our first hypothesis:

Hypothesis 1. *In situations with conflicting interests under full information, the share of self-serving decisions is lower in the RECOMMEND treatment than in the SELL treatment (moral engagement).*

Note that, up to this point, Hypothesis 1 is also consistent with models based on additively separable cost of lying, which can only occur in the RECOMMEND treatment. Such cost would increase the left hand-side of condition (4), thus having a similar effect under FULLINFO as differences in norm activation. Note, however, that for the prediction to hold qualitatively, introducing such cost is not necessary, as even for $\eta_s = \eta_r$, condition (4') is already laxer than condition (1) if buyers' trust is expected to be incomplete.

4.3.2 Hidden but revealable Information

Let us now introduce the possibility of information avoidance as an additional strategy to reduce cognitive dissonance, as in our HIDDENINFO treatments. In this case, the true value of v is initially unobservable, which implies that it is *a priori* unclear if interests are aligned or conflicting. Denote with μ the *ex ante* probability of interests being aligned, i.e. $\mu = P(v = \bar{v})$, which is 0.5 in our experiment. Furthermore, we use index $k \in \{0, 1\}$ to denote the state of the individual's level of information, with $k = 1$ representing a situation where the information is revealed, while $k = 0$ represents non-revelation.

For an uninformed individual ($k = 0$) for which (4) holds in a RECOMMEND situation, the expected costs of cognitive dissonance Φ_0 when choosing option $\neg a$ are determined by the individual's (subjective) beliefs on the probability of aligned interests as

follows:

$$\Phi_0^r = \Phi \left(\eta_r - \hat{E}(v); \alpha \right)$$

Notice that, even without subjective distortions in beliefs, Φ_0 is always smaller than under certainty, as represented in (4), which creates an incentive to simply remain uninformed and choose the self-serving option. This strategy, however, represents a sort of self-deception which, in turn, is likely to be associated with a feeling of displeasure with one's own self-serving rationalization. To take this into account, we again follow [Rabin \(1994\)](#) and the subsequent literature by introducing costs of self-deception Ψ_k dependent on the existence of an (activated) norm which become smaller with the amount of available information revealed, i.e. the value of k . Hence, the costs of self-deception, for k signals revealed, are

$$\Psi_k = \Psi_k((\hat{\mu}_k - \mu_k), \eta).$$

Notice that such costs are also increasing in the misperception of probability μ , which is of lesser interest here, as $\mu = 1/2$ in our experiment, which is usually not subject to probability weighting. We will therefore omit this term for ease of notation. Given these specifications and for k signals revealed, the valuation of the self-serving option $\neg a$ can be written as

$$(5) \quad EU_k(\neg a) = E[p \mid \neg a] - \Phi_k \left(\eta_r - \hat{E}(v \mid k) \right) - \Psi_k(\eta_r).$$

Note that for the signal revealing the truth with certainty, as assumed here, there is no self-deception if the signal is revealed. Hence, for $k = 1$, equation (5) reduces to (2), as $\Psi_1 = 0$. Furthermore, we assume that no cognitive dissonance is felt if the outcome is unknown, i.e. $\Phi_0 = 0$. This assumption, while plausible, is made for ease of notation, and does not alter the qualitative results of what follows.

To identify the potential for self-serving information avoidance, consider an altruistic seller for whom, under FULLINFORMATION, (4) holds, i.e. under FULLINFORMATION they would want to choose the altruistic option. For simplicity, we assume risk-neutrality. This individual's expected valuation when planning to reveal the information, but before doing so, is

$$(6) \quad EU_{k=1}^r = \mu \cdot E[p \mid a] + (1 - \mu) \cdot E[p \mid \neg a].$$

In this case, given that we assume (4) to hold, the seller will only choose the self-serving option if interests are aligned, otherwise option $\neg a$ will be recommended.

As the individual intends to always choose the option corresponding to the personal norm, no cost from cognitive dissonance will arise.

For such an individual, the decision to remain uninformed is determined via a comparison of (5) with $k=0$ and (6). More precisely, the seller will only reveal the information and choose the altruistic option iff $EU_0(-a) < EU_{k=1}$ which is the case for

$$(7) \quad \Phi(\eta_r - \underline{v}) + z > \Delta_p \wedge \frac{\Psi_0(\eta_r)}{\mu} + z > \Delta_p.$$

The first condition in (7) corresponds to (4'). Thus, if the first condition does not hold, the seller would choose opportunism even under full information. If the first condition does hold, but the second does not, the seller will avoid the information and choose the option with the higher payoff to himself (i.e. the opportunistic option in a conflicting interest situation). Information will be avoided if, in the uninformed state, the costs of self-deception, which depend on norm activation, are not too high compared to the difference in payoffs.

Thus, for the RECOMMEND condition there exists a range of situations where sellers with a personal norm would choose the altruistic recommendation in the FULLINFO condition, but not under HIDDENINFO. This phenomenon represents the exploitation of moral wiggle room via self-serving avoidance of information. Following the literature on moral wiggle room, this type of behavior is observed if the following hypothesis is confirmed:

Hypothesis 2. *In situations with conflicting interests in the RECOMMEND treatments, the share of selfish choices is lower in the FULLINFO condition than in the HIDDENINFO condition.*

Deriving the condition for self-serving information avoidance for SELL situations is analogous to the considerations above. In this case, the seller will choose the altruistic option iff

$$(8) \quad \Phi(\eta_s - \underline{v}) > \Delta_p \wedge \frac{\Psi_0(\eta_s)}{\mu} > \Delta_p.$$

Thus, for those sellers with an activated norm, we can expect a similar behavior as in the RECOMMEND case, i.e. self-serving information avoidance if the first condition in (8) does hold, but the second does not. This leads to the following hypothesis:

Hypothesis 3. *In situations with conflicting interests in the SELL treatments, the share of selfish choices is lower in the FULLINFO condition than in the HIDDENINFO condition.*

The latter hypothesis, however, has a lower probability of being confirmed than the one for the RECOMMEND case, in particular if norm activation is significantly lower in the SELL treatments. For all subjects for whom the norm is not activated, the left-hand sides of both conditions in (8) reduce to 0, in which case the seller will choose the opportunistic option.

If we compare only those sellers with an activated norm, we notice that the second condition in (8) is stricter for SELL situations than in (7) for RECOMMEND situations. Recall that the sellers, for whom the lower boundary is binding, will always reveal information *and* choose the altruistic option. They would hence behave differently to ‘curious’ egoists, who will also reveal information, but choose the self-serving option. Thus, if exploitation of moral wiggle room is more frequent for the RECOMMEND treatments, the share of altruistic choices for decisions with revealed information under HIDDENINFO should be larger for the RECOMMEND than for the SELL treatment.

Obviously, Hypotheses 2 and 3 would also hold qualitatively if we added a situation-independent cost of lying. Such cost would only affect the left-hand side of conditions (7) and (8), while the right-hand side remained the same for both, RECOMMEND and SELL situations. To see this, notice that costs of lying neither enter expected utility (6) for $k = 0$, nor expected utility (5), as under uncertainty there is no intentional lie. Such costs would, however, affect choices made *after* information is revealed. In the literature, costs of lying are assumed to depend on the difference in both payoffs, which corresponds to the variable pairs WeakCI/StrongCI/Symmetric and HighPD/LowPD in our parametrization. Thus, if such costs of lying were driving the sellers’ behavior, the parametrization should have an effect on decisions after revelation under the RECOMMEND condition compared to SELL.

In contrast, within the model above, the parametrization does not affect decisions *after* information is revealed. If, for any pairs of WeakCI/StrongCI/Symmetric and HighPD/LowPD, condition (7) holds, subjects who reveal will choose the altruistic option. Otherwise, they choose opportunistically. Thus, if we are confident that additively separable costs of lying do not play a significant role here, we expect our last hypothesis to be rejected:

Hypothesis 4. *The effects of the parametrization on the choices of willingly informed sellers differ across treatment variations (SELL vs. RECOMMEND).*

5 Results

To get an overview of the main effects revealed within the experiment, we first analyze the aggregated data. In a second step, we study the treatment effects which allows us to test our hypotheses and confirms our preliminary results. Subsequently, we proceed

with an analysis of the parametrization to further substantiate the behavioral effects at play. As we are interested in situations with a potential moral conflict, we exclusively consider the share of selfish actions in situations with conflicting interests – i.e. situations where the option maximizing the seller’s payoff minimizes the buyer’s payoff and *vice versa*.

5.1 Analysis of Aggregated Data

We observe that sellers give selfish recommendations in 54.5% of all cases in the RECOMMEND treatments, while they sell the product that maximizes their payoff and minimizes their client’s payoff in 64.1% of all situations in the SELL treatments. This difference in the share of opportunistic decisions is highly significant at a p-value smaller than 0.01 in a χ^2 -test that allows for clustering at the seller level. Thus, it seems that sellers are more likely to take the buyer’s needs into account when the buyer is an active decision-maker. Note that this difference cannot be explained on the grounds of simple models with altruistic types, as (for sellers with stable altruistic preferences) perfect separation of types would arise in both treatments. Hence, the simple fact that the buyer has a choice seems to reduce the sellers’ propensity to act opportunistically.

Yet, at this point in the analysis, we cannot clearly attribute this treatment effect to a behavioral cause. As the share of selfish recommendations in the RECOMMEND treatments is 54.5%, which is relatively close to 50%, it is also possible that sellers simply randomize in this treatment. This would correspond to the prediction with completely self-interested, rational agents on both market sides, as discussed in Section 4.1. However, with the subsequent, more detailed analyses, this explanation is to be rejected, as we find that the parametrization influences recommendations significantly (see Table 5).

For the aggregated data, we observe that 54.3% of the sellers’ choices are opportunistic under FULLINFORMATION and 65.2% are selfish in the HIDDENINFORMATION treatments. Again, the difference is highly significant ($p < 0.01$) in the clustered version of the χ^2 -test. As information was revealable without cost, this result is in line with the prediction of sellers exploiting moral wiggle room via information avoidance.

To substantiate the results of the χ^2 -tests, we perform regression analyses (see Table 2) to control for the parametrization of the respective decision situations as well as potential time trends. In random-effects panel regressions, we regress an indicator variable that takes the value of one if the decision was selfish on an indicator variable capturing the treatment variation – RECOMMEND vs. SELL and HIDDENINFORMATION vs. FULLINFORMATION. We further include indicator variables for decision situations with a weak conflict of interests, symmetric decision situations and a high price differ-

ence for the seller. In addition, we control for potential time trends by including the decision number, as well as for the subject’s age and self-identified gender.

Table 2: Regressions: Share of selfish choices on an aggregate basis

	Recommend vs. Sell	Hidden vs. Full
Recommend	-0.090*** (0.029)	
Hidden		0.088*** (0.029)
PDhigh	0.073*** (0.015)	0.074*** (0.015)
WeakCI	0.479*** (0.024)	0.475*** (0.024)
Symmetric	0.043** (0.021)	0.039* (0.021)
DecisionNo	0.003 (0.003)	0.002 (0.003)
Age	-0.002 (0.001)	-0.001 (0.001)
Male	0.020 (0.029)	0.026 (0.029)
Constant	0.424*** (0.056)	0.338*** (0.059)
R ²	0.254	0.250
N	1834	1834

Notes: Output from random-effects panel regressions. Standard errors (in parentheses) are clustered on subject-level. The dependent variable is a dummy which takes the value of 1 if the selfish option is chosen in a conflict situation. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The regressions confirm our previous results: decisions are more selfish in the SELL treatments as opposed to the RECOMMEND treatments.⁶ Sellers’ decisions are also more opportunistic under HIDDENINFORMATION than in the FULLINFORMATION treatments. We further observe that decisions are more opportunistic if the price difference is high and when the seller faces a symmetric decision situation or a weak conflict of interests as opposed to a strong conflict of interests, with the latter effect being more pronounced. This is consistent with the model presented in the previous section. Neither experience nor the seller’s age or gender have a significant effect on opportunistic choices.

5.2 Treatment Effects

To investigate the influence of each treatment variation on seller behavior, we again perform a random-effects panel regression (Table 3) including the same control vari-

⁶Note that our hypotheses derived in the previous section are tested by use of the regression results reported in Table 3, which were also tested for multiplicity.

ables as before, but adding an interaction term of the treatment variation with respect to information and recommendation. The marginal effects are presented in the bottom part of the table.

Table 3: Regression and marginal effects: Share of selfish choices

	Selfish choices
Hidden	0.055 (0.041)
Recommend	-0.120*** (0.040)
Hidden*Recommend	0.065 (0.059)
WeakCI	0.476*** (0.024)
Symmetric	0.039* (0.021)
PDhigh	0.074*** (0.015)
Age	-0.001 (0.001)
Male	0.024 (0.029)
DecisionNo	0.002 (0.003)
R ²	0.260
N	1834
Marginal Effects	
Recommend	
if Full = 1	-0.120*** (0.040)
if Hidden = 1	-0.055 (0.041)
Hidden	
if Sell = 1	0.055 (0.041)
if Recommend = 1	0.120*** (0.041)

Notes: Output from random-effects panel regressions. Standard errors (in parentheses) are clustered on subject-level. The dependent variable is a dummy which takes the value of 1 if the selfish option is chosen in a conflict situation. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. To test for multiplicity, we computed sharpened False Discovery Rate (FDR) q-values using the method described in [Anderson \(2008\)](#). Using this correction, the significance levels remain the same as reported above.

To discuss these results, we present the relations between treatments in Table 4, which shows the share of opportunistic choices for each treatment and the effects identified in our treatment comparison. We observe that the share of selfish choices is lowest in the RECOMMEND treatment under FULLINFORMATION, while it is highest in the

SELL treatment under HIDDENINFORMATION. We summarize our observations from Tables 3 and 4 as follows:

Result 1. *The treatment variations influence seller behavior as follows:*

- (i) *Under FULLINFORMATION, sellers' opportunistic decisions are less frequent if sellers make recommendations instead of sales decisions (moral engagement).*
- (ii) *Opportunistic recommendations are more common when sellers need to perform a diagnosis than when they directly observe the buyer's needs (moral wiggle room).*
- (iii) *The necessity to perform a diagnosis does not influence the share of opportunistic decisions of direct sales decisions.*
- (iv) *Under HIDDENINFORMATION, recommendations and sales decisions are equally selfish.*

Under FULLINFORMATION, the sellers' propensity to opportunistic behavior is significantly smaller if the buyer is left with the choice not to follow the recommendation. This cannot be explained by a model with binary types and stable preferences, but is consistent with our cognitive dissonance model, as Result 1 (i) confirms our Hypothesis 1. The result is also a first indication of potential moral disengagement, but from this result alone it cannot yet be inferred that there are differences in norm activation between both treatments.

We also find a treatment effect for the RECOMMEND treatments that is consistent with the exploitation of moral wiggle room via information avoidance, i.e. our Hypothesis 2 is confirmed. Interestingly, this effect cannot be detected in a comparison of the SELL treatments (Result 1 (iii)). Thus, our Hypothesis 3 is to be rejected. As mentioned in our discussion of this hypothesis, this again hints to lower levels of norm activation under SELL.

Table 4: Treatment effects: Share of selfish choices

		OBSERVABILITY OF THE BUYER'S NEEDS				
		FULLINFO		HIDDENINFO		
ACTION	RECOMMEND	0.482	⇒	<i>moral wiggle room</i>	⇒	0.621
		↓		<i>(Result 1 (ii))</i>		⇕
		<i>moral disengagement</i>				<i>no sign. effect</i>
		<i>(Result 1 (i))</i>				<i>(Result 1 (iv))</i>
		↓				⇕
	SELL	0.606	⇔	<i>no sign. effect</i>	⇔	0.680
				<i>(Result 1 (iii))</i>		

Note that the effects identified above are significant while controlling for the parametrization. In particular, we control for the relative intensity of the conflict of interest,

via the use of variables ‘WeakCI’, indicating if a situation was associated with a weak conflict of interests, and ‘Symmetric’ indicating a symmetric decision situation. As shown in Table 1, WeakCI is the only type of situation for which Kaldor-Hicks efficiency does not increase with decreasing opportunism. Thus, for all other choice situations, our results on altruistic behavior also carry over to considerations on market efficiency. We conclude that for the information asymmetry investigated here, both effects presented in Table 4 lead to a decrease in efficiency of our decision situations where opportunism is Kaldor-Hicks inefficient. Still, it is remarkable that the share of opportunistic choices never exceeds 68%. Thus, altruistic choices on the part of the seller are significantly more frequent than standard microeconomic models on information asymmetries in markets predicts.

5.3 The Role of the Parametrization

The conclusions drawn above are confirmed when we analyze the impact of the parametrization on selfish decisions directly (see Figure 1). We observe that, in all treatments, sellers indeed react to the intensity of the conflict of interests as our model predicts: opportunistic decisions are more frequent if the conflict of interests is weak and less frequent with a strong conflict of interests. The share of selfish choices in symmetric decision situations is comparable to the one in strong conflict situations. While choices are more opportunistic in case of a higher price difference when sellers face a strong conflict of interests, the price difference does not affect opportunism in situations with a weak conflict of interests. As predicted by our model, the level of opportunism varies across treatments, while the influence of the parametrization is similar. The regressions in Table 5 confirm this result: decisions are more opportunistic when the selfish option ensures a relatively high payoff for the buyer (“Weak CI”).

The regressions in Table 5 confirm this result: decisions are more opportunistic when the selfish option ensures a relatively high payoff for the buyer (“Weak CI”).

We can also confirm the above-made considerations on efficiency. Recall that for all parametrizations except “WeakCI”, the choice that maximizes the buyer’s payoff coincides with the choice maximizing market efficiency. At first sight, one might take this as an indication that sellers also act according to a preference for efficiency. Yet, this does not seem to be the case when we consider the results of both choice situations that fall under “WeakCI”. In fact, we observe no significant difference in selfish decisions between these situations. Recall from Table 1 that “WeakCI HighPD” is the only situation where a reduction in opportunism decreases efficiency, while for “WeakCI LowPD”, Kaldor-Hicks efficiency remains the same for both options. Yet, if sellers’ choices were driven by efficiency concerns, choices for both types of situations should

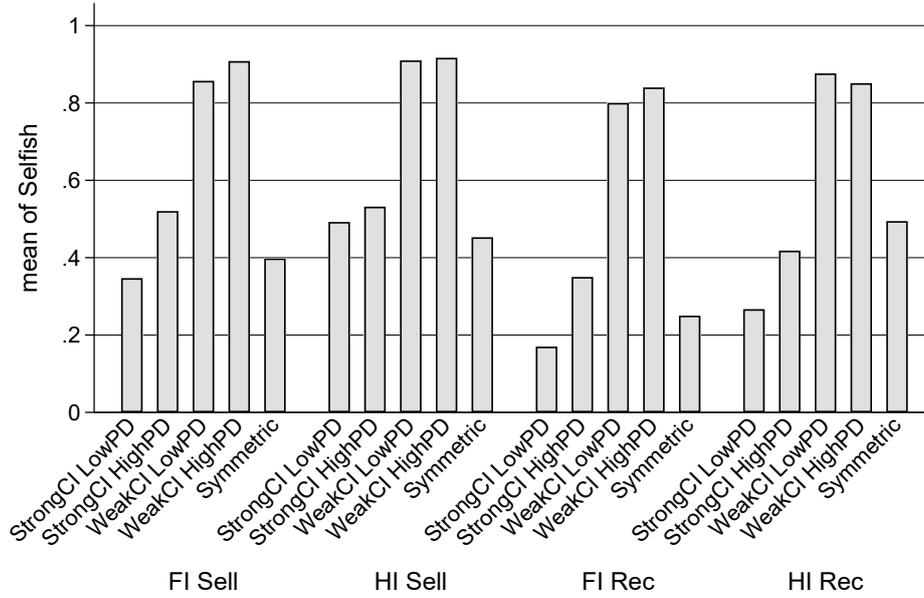


Figure 1: Share of selfish choices for each parametrization across treatments

Table 5: Regressions: Selfish choices in each treatment

	(1)	(2)	(3)	(4)
	FI Sell	HI Sell	FI Rec	HI Rec
WeakCI	0.466*** (0.051)	0.394*** (0.052)	0.563*** (0.048)	0.489*** (0.054)
Symmetric	0.003 (0.043)	-0.062 (0.049)	0.044 (0.040)	0.133*** (0.049)
PDhigh	0.125*** (0.033)	0.023 (0.026)	0.105** (0.041)	0.052 (0.040)
DecisionNo	-0.006 (0.009)	-0.000 (0.007)	0.001 (0.006)	0.005 (0.008)
Age	-0.000 (0.003)	0.000 (0.002)	-0.002 (0.002)	-0.003 (0.003)
Male	0.050 (0.061)	0.036 (0.056)	0.039 (0.053)	-0.037 (0.061)
Constant	0.392*** (0.123)	0.479*** (0.095)	0.236** (0.093)	0.437*** (0.124)
R ²	0.230	0.209	0.318	0.243
N	490	435	500	409

Notes: Output from random-effects panel regressions. Standard errors (in parentheses) are clustered on subject-level. Dependent variable is a dummy which takes the value of 1 if the selfish option is chosen in a conflict situation. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

differ significantly. As this is clearly not the case here, we may conclude that efficiency concerns do not drive our results.

5.4 Diagnosis

The parametrization also influences the sellers' decision to reveal information. As shown in Figure 2, sellers are less likely to gather information in situations with a weak conflict of interests than with a strong conflict. Again, this is consistent with the prediction of our cognitive dissonance model, as in case of strong conflicts of interest the associated costs of self-deception become too large for information avoidance to be a viable option to resolve dissonance. For situations with a strong conflict of interests, their willingness to reveal information is slightly reduced when the price difference is high, as opposed to situations with a low price difference. Yet, in a comparison of treatments for the aggregated data, at first sight, there does not seem to be a significant treatment effect in information revelation: We find that in 54.5% of the situations, sellers perform a diagnosis in the SELL treatment. In the RECOMMEND treatment this share is lower with sellers performing a diagnosis in only 48.8% of the situations. Yet, with a p-value of 0.285 in a χ^2 -test adjusted for clustering, this difference is not significant.

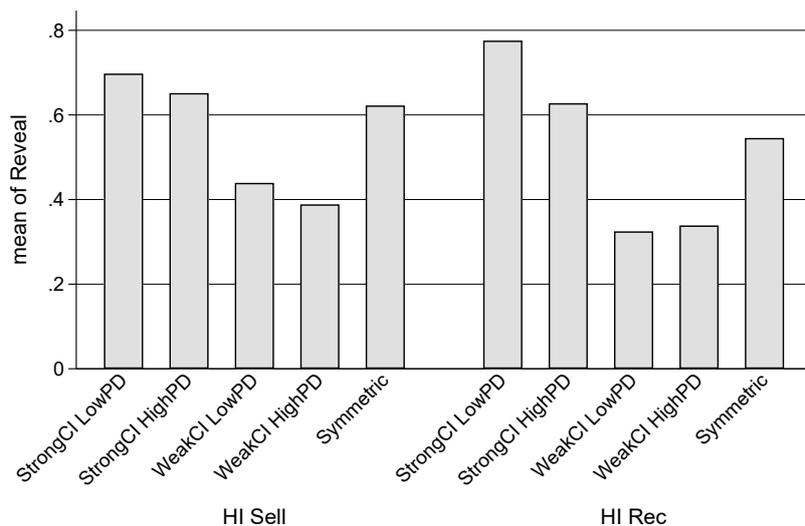


Figure 2: Share of sellers revealing the state of the world under different parameter settings

However, a comparison of the revelation rate is not sufficient to draw any conclusion on the tendency to exploit moral wiggle room as information is not only revealed for higher levels of norm conformity (as predicted via conditions (7) and (8)). In addition, we have to account for potential 'curious egoists', who tend to opportunism even if perfectly informed. Such subjects perform the costless diagnosis, even if this does not affect their (opportunistic) choice.

Thus, to establish that exploitation of moral wiggle room can account for the differences in both HIDDENINFO treatments, we need to compare the shares of opportunistic decisions for those decisions that were taken after information was revealed. These are

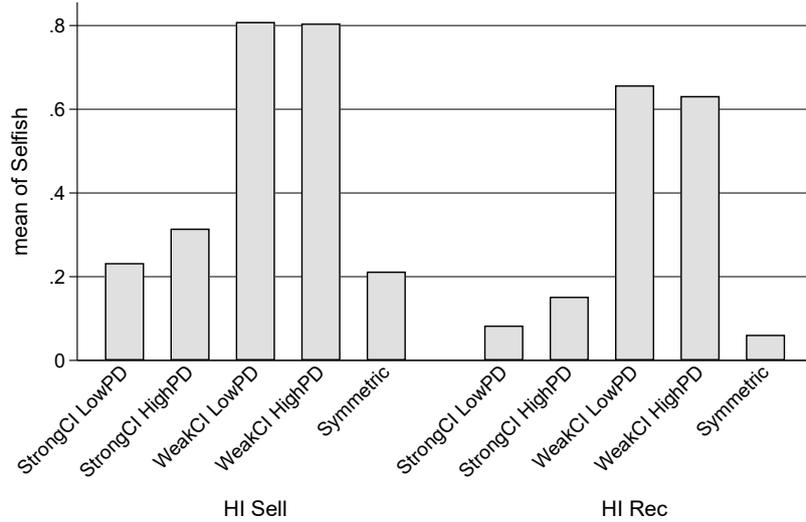


Figure 3: Share of selfish choices for each parametrization of willingly informed sellers

presented in Figure 3, which indicates that there is indeed a difference between SELL and RECOMMEND treatments in the size of these shares, while the overall pattern of parameter-induced effects remains similar. Again, we confirm that there is indeed a significant treatment effect via a regression, represented in Table 6. This indicates that norm activation is indeed too low in the SELL treatment to find a measurable amount of exploitation of moral wiggle room.

Notice also from the regression in Table 6 that there is no indication that the differences in opportunism between SELL and RECOMMEND could be explained via additively separable costs of lying. In the literature, such costs are dependent on the payoff of the buyer, the seller, and most often both. Yet, the respective interaction terms do not indicate significant effects. Obviously, the insignificance of these effects is not altered if we test for multiplicity. Similarly, marginal effects of the parametrization (reported at the bottom of the table) are not systematically different for SELL and RECOMMEND. Hypothesis 4 is hence to be rejected, and we can state as a result:

Result 2. *The impact of the parametrization on the choices of willingly informed sellers does not differ across treatment variations (SELL vs. RECOMMEND). Additively separable costs of lying are hence not confirmed as an explanation for the treatment effects.*

Thus, a large part of the patterns observed here are consistent with the cognitive dissonance model derived in the previous section. In contrast to additively separable costs of lying, differences in norm activation remain plausible as a mechanism reinforcing or dampening cognitive dissonance in the RECOMMEND and SELL treatments.

Table 6: Regression: Selfish choices of willingly informed subjects

	Selfish Choices
Recommend	-0.162** (0.066)
PDhigh	0.052 (0.049)
Rec*PDhigh	-0.021 (0.070)
WeakCI	0.517*** (0.070)
Symmetric	-0.060 (0.059)
Rec*WeakCI	0.003 (0.103)
Rec*Symmetric	0.031 (0.073)
Age	-0.001 (0.002)
Male	0.049 (0.047)
DecisionNo	-0.002 (0.007)
Constant	0.277*** (0.093)
R ²	0.328
N	467
Marginal Effects	
PDHigh	
if Recommend = 1	0.031 (0.051)
WeakCI	
if Recommend = 1	0.520*** (0.076)
Symmetric	
if Recommend = 1	-0.029 (0.041)

Notes: Output from random-effects panel regressions. Standard errors (in parentheses) are clustered on subject-level. Dependent variable is a dummy which takes the value of 1 if the selfish option is chosen in a conflict situation. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.5 Buyers' Decisions and Overtreatment

While we are mainly interested in the sellers' behavior, some interesting results can be derived from the subsequent decision of the buyers in the RECOMMEND treatments. Here, we observe that buyers are more likely to follow the recommendation in the HIDDENINFORMATION treatment (78.3%) than in the FULLINFORMATION treatment (71.9%) – a difference that is significant at a p-value of 0.012 in a χ^2 -test adjusted for

clustering. This result is particularly surprising as buyers do not know if their matched seller decided to perform a diagnosis or decided to remain uninformed. Hence, buyers do not seem to expect sellers to exploit moral wiggle room.

Figure 4 shows the buyers' decisions in more detail. When confronted with a recommendation that does *not* maximize the seller's payoff (left panel), the share of buyers following the recommendation varies significantly with the parameterization. This is at odds with the predicted behavior of perfectly rational agents, as such a recommendation can be inferred to always be truthful.⁷

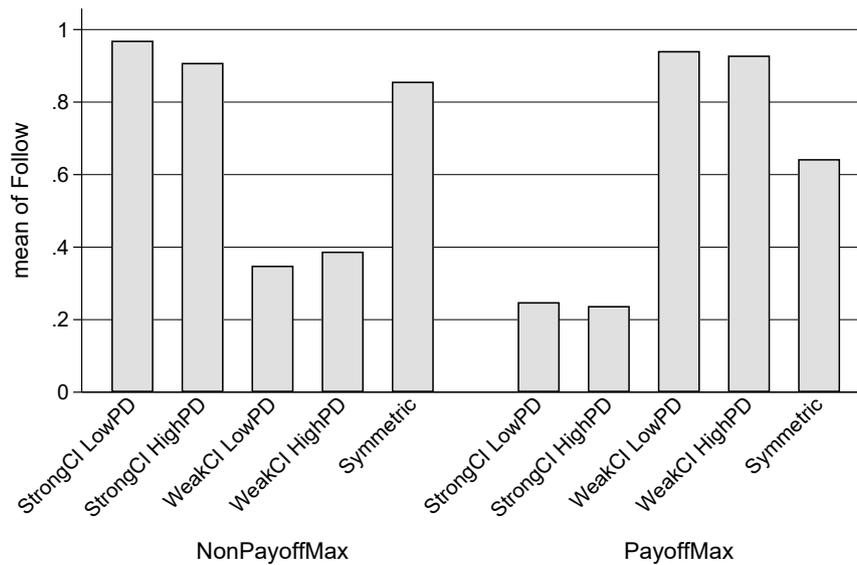


Figure 4: Share of buyers following the recommendations under different parameter settings

The observation that buyers seem to show a high level of trust in an obviously non-opportunistic recommendation in situations with a strong conflict of interests (left panel of Figure 4) might instead be driven by buyers choosing the 'safer' option. Recall that with a strong conflict of interest, following the non-opportunistic recommendation coincides with choosing the option that guarantees the same, and relatively high, monetary payoff in both states of the world. A tendency to choose the 'safer' option can also explain buyer's behavior for such a recommendation with a weak conflict of interests. In this case, they are more likely to choose the non-recommended option, which is associated with lower risk in case the recommendation is not truthful. In symmetric situations, buyers would be indifferent between both options if they interpreted the seller's recommendation as an empty signal. The observation that they tend to follow the recommendation may hence be interpreted as buyers exhibiting a positive amount of trust in the recommendation.

⁷This holds at least if sellers cannot be anticipated to be spiteful, which is implausible in a one-shot setting with perfect anonymity.

As to buyers' reaction to a recommendation that maximizes the seller's payoff (right panel of Figure 4), we again observe that buyers exhibit a strong preference for the safe option: in situations with a strong conflict of interests, buyers prefer not to follow the seller's advice. In contrast, in a situation with a weak conflict of interests, buyers have a stronger tendency to follow the sellers' advice, which coincides with the safer option.

In both panels of Figure 4, we can see that buyers do not seem to differentiate between situations "WeakCI HighPD" and "WeakCI LowPD". Again, this is at odds with standard information economics, according to which the seller's opportunism is mainly determined by his own payoffs. Buyers in our sample do not seem to use this insight to infer that under "WeakCI HighPD" average recommendations should be more selfish than under "WeakCI LowPD". Instead, the majority of buyers seem to follow the simple heuristic of choosing the option with the lowest possible losses.

Hence, in the terminology of the credence goods literature, we find that buyers self-select into (potential) overtreatment, i.e. they choose the option that guarantees a relatively high level of payoff in both states of the world. This would be in line with choosing a treatment that solves both the minor and the major problem, albeit at a larger price than the solution that only addresses the minor problem.

Result 3. *Buyers exhibit a preference for the safer option and hence select into overtreatment.*

Note, however, that our setup precluded that sellers could adjust to these tendencies in buyer behavior, as the results of all sales decisions were only communicated at the end of the experiment.

6 Conclusion

In this experiment, we analyze the effects of two different design choices in credence good experiments on the moral behavior of the seller. First, we investigate the effect of the buyer being able not to follow the seller's recommendation versus an automatic sale of the seller's chosen treatment. As laid out in the introduction, both varieties exist in real-world contracts, involving, for example, car repairs, IT-services, or any type of craftsmanship. The second variation analyzed here is the necessity for sellers to perform a diagnosis. As we are exclusively interested in moral motives in the seller's decision to perform a diagnosis, information costs are 0. We compare a treatment where diagnosis is necessary to a treatment with full information where the seller does not need to reveal any information in order to identify the product that best suits the buyers' needs. We implement these treatments using a 2x2 factorial design, while the parametrization of the decision situation is varied within-subject via 10 independent transactions.

We observe that the level of the sellers' opportunistic choices is significantly lower if they are to recommend an option to the buyer than if they are able to sell the product without additional buyer intervention. This is in contrast to the predictions based on a standard-microeconomic individual, and leads to a lower level of opportunism than standard theory would predict, as the exploitation of asymmetric information on the part of the seller is mitigated by their (situation-dependent) moral preferences. Without buyer intervention, the level of opportunism is closer to the standard prediction, while sellers' choices are still not perfectly self-serving. In this context, it is important to note that a reduction in opportunism leads to an increase in Kaldor-Hicks efficiency in all but two of our decision situations.

As to the revelation of information, we find that the seller tends to exploit moral wiggle room by avoiding costless information. Thus, a significant part of the sellers remain uninformed with respect to the buyer's needs in order to sell the self-serving option without having to incur potential psychological costs from a bad conscience. Again, this is in contradiction with predictions based on a neoclassical individual, which would reveal information if it was without cost. While the exploitation of moral wiggle room has been observed in a variety of contexts (see, e.g., [Dana et al., 2007](#); [Grossman and van der Weele, 2017](#); [Conrads and Irlenbusch, 2013](#); [Momsen and Ohndorf, 2020b](#)), this is the first experiment where this phenomenon was observed within the behavior of sellers.

Interestingly, if buyers cannot choose the option that is ultimately sold, the necessity to perform a diagnosis does not further influence the sellers' choices, as in these cases, the level of opportunism is already relatively high even under perfect information. This lends support to the idea that moral norms need to be activated to influence behavior ([Schwartz, 1977](#)). The activation of a moral norm seems more likely if the buyer is perceived as a human-being with a choice than otherwise. While the alternative explanation of additively separable cost of lying could explain the treatment difference under full information, it is not consistent with the entirety of the investigated cases.

Moreover, we observe that the influence of our parametrization is consistent with our model: opportunistic decisions are more frequent when the difference in the buyer's payoff between the two options is relatively small. Hence, when mistreating the buyer has only mild negative consequences, sellers are willing to do so to maximize their own payoff. This hints to the fact that, in the terminology of credence goods, overtreatment is the prevalent problem of opportunistic behavior in such situations. This conclusion holds at least if we assume that buyers' negative consequences are larger with undertreatment, i.e. the provision of a product or service that is not suitable, such that the buyer's problem persists after the purchase. While we do not consider overcharging

per se, our results suggest that it might be frequent, as the buyer's needs are fulfilled without causing major harm to the buyer.

The parametrization also has a major impact on buyers' decisions on whether or not to follow the seller's advice. Most importantly, buyers exhibit a strong preference for the safe option and thus prefer not to follow the advice when the seller recommends the riskier option. Interpreted in the context of credence goods, this suggests that buyers have a tendency for self-selecting into being overtreated as this option definitely fulfills their needs. As sellers also seem to show a tendency to engage in overtreatment, our results suggest that overtreatment is likely to be the prevalent type of opportunism in the type of credence goods markets considered here.⁸

The fact that the market outcome varies significantly between our treatments indicates that the specific design of the credence goods market has an influence on the level of opportunism that can be observed. As, in principle, we implemented different variations that are common in the experimental literature on credence goods, the design choices for these markets do not seem innocent. This would be particularly important in those cases, where the reduction in opportunism increases the efficiency of the market.

More tentatively, our results also lend themselves to inform the general discussion on markets and morals, as it shows that even minor alterations in the characteristics of the analyzed market can have an impact on the outcome. Thus, when debating morals within markets, it seems prudent to qualify any general conclusions, and clearly state the specificities of the market in question. In particular, arguments that explain moral erosion with the anonymity of market participants would have to be qualified. As our results suggest, the activation of moral norms is still possible in situations of relatively high anonymity. What seems to be important is the perception of the opposite market side as a human-being capable to make decisions on their own.

⁸Note, however, that this conclusion is only valid as long as the cost of overtreatment is not too large.

References

- Alger, I. and Salanié, F. (2006). A Theory of Fraud and Overtreatment in Experts Markets. *Journal of Economics & Management Strategy*, 15(4):853–881.
- Anagol, S., Cole, S., and Sarkar, S. (2017). Understanding the Advice of Commissions-Motivated Agents: Evidence from the Indian Life Insurance Market. *The Review of Economics and Statistics*, 99(1):1–15.
- Anderson, M. L. (2008). Multiple inference and gender differences in the effects of early intervention: A reevaluation of the abecedarian, perry preschool, and early training projects. *Journal of the American Statistical Association*, 103(484):1481–1495.
- Bae, M. (2016). The effects of anonymity on computer-mediated communication: The case of independent versus interdependent self-construal influence. *Computers in Human Behavior*, 55:300–309.
- 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. *The Review of Economic Studies*, 80(3):876–891.
- Balafoutas, L. and Kerschbamer, R. (2020). Credence goods in the literature: What the past fifteen years have taught us about fraud, incentives, and the role of institutions. *Journal of Behavioral and Experimental Finance*, 26:100285.
- Balafoutas, L., Kerschbamer, R., and Sutter, M. (2017). Second-degree moral hazard in a real-world credence goods market. *The Economic Journal*, 127(599):1–18.
- Bandura, A. (2002). Selective moral disengagement in the exercise of moral agency. *Journal of moral education*, 31(2):101–119.
- Bandura, A., Underwood, B., and Fromson, M. E. (1975). Disinhibition of aggression through diffusion of responsibility and dehumanization of victims. *Journal of Research in Personality*, 9(4):253–269.
- Bartling, B., Weber, R. A., and Yao, L. (2015). Do markets erode social responsibility? *Quarterly Journal of Economics*, 130(1):219–266.
- Beck, A., Kerschbamer, R., Qiu, J., and Sutter, M. (2014). Car mechanics in the lab—Investigating the behavior of real experts on experimental markets for credence goods. *Journal of Economic Behavior & Organization*, 108:166–173.
- Bénabou, R. and Tirole, J. (2016). Mindful economics: The production, consumption, and value of beliefs. *Journal of Economic Perspectives*, 30(3):141–64.

- Blasch, J. and Ohndorf, M. (2015). Altruism, moral norms and social approval: Joint determinants of individual offset behavior. *Ecological Economics*, 116:251–260.
- Bolton, P. and Dewatripont, M. (2005). *Contract Theory*. MIT press.
- Califano, G., Capasso, M., and Caso, D. (2022). Exploring the roles of online moral disengagement, body esteem, and psychosexual variables in predicting sexting motivations and behaviours. *Computers in Human Behavior*, 129:107146.
- Chen, D. L., Schonger, M., and Wickens, C. (2016). otree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9:88–97.
- Cherry, T. L., Frykblom, P., and Shogren, J. F. (2002). Hardnose the dictator. *American Economic Review*, 92(4):1218–1221.
- Conrads, J. and Irlenbusch, B. (2013). Strategic ignorance in ultimatum bargaining. *Journal of Economic Behavior & Organization*, 92:104–115.
- Dana, J., Weber Jason Xi Kuang, R. A., Mel-Ion, C., Dana, J., Weber, R. A., and Kuang, J. X. (2007). Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. *Economic Theory*, 33:67–80.
- Darby, M. R. and Karni, E. (1973). Free Competition and the Optimal Amount of Fraud. *The Journal of Law and Economics*, 16(1):67–88.
- De Groot, J. I. and Steg, L. (2009). Morality and prosocial behavior: The role of awareness, responsibility, and norms in the norm activation model. *The Journal of Social Psychology*, 149(4):425–449.
- De Jaegher, K. (2012). The Value of Private Patient Information in the Physician-Patient Relationship: A Game-Theoretic Account. *Computational and Mathematical Methods in Medicine*, 2012:16.
- Dellavigna, S., List, J. A., Malmendier, U., and Rao, G. (2017). Voting to tell others. *The Review of Economic Studies*, 84:143–181.
- Dulleck, U. and Kerschbamer, R. (2006). On Doctors, Mechanics, and Computer Specialists: The Economics of Credence Goods. *Journal of Economic Literature*, 44(1):5–42.
- Dulleck, U. and Kerschbamer, R. (2009). Experts vs. discounters: Consumer free-riding and experts withholding advice in markets for credence goods. *International Journal of Industrial Organization*, 27:15–23.

- Dulleck, U., Kerschbamer, R., Sutter, M., Beck, A., Dittrich, D., Emons, W., Kroll, S., Luhan, W., and Samuelson, L. (2011). The Economics of Credence Goods: An Experiment on the Role of Liability, Verifiability, Reputation, and Competition. *American Economic Review*, 101:526–555.
- Emons, W. (1997). Credence Goods and Fraudulent Experts. *The RAND Journal of Economics*, 28(1):107–119.
- Emons, W. (2001). Credence goods monopolists. *International Journal of Industrial Organization*, 19:375–389.
- Epley, N. and Gilovich, T. (2016). The mechanics of motivated reasoning. *Journal of Economic perspectives*, 30(3):133–40.
- Exley, C. L. and Kessler, J. B. (2021). Information avoidance and image concerns. Technical report, National Bureau of Economic Research.
- Falk, A. and Szech, N. (2013). Morals and markets. *Science*, 340(6133):707–711.
- Fong, Y.-F. (2005). When Do Experts Cheat and Whom Do They Target? *The RAND Journal of Economics*, 36(1):113–130.
- Fong, Y. F., Liu, T., and Wright, D. J. (2014). On the role of verifiability and commitment in credence goods markets. *International Journal of Industrial Organization*, 37:118–129.
- Gneezy, U. (2005). Deception: The role of consequences. *American Economic Review*, 95(1):384–394.
- Gneezy, U., Rockenbach, B., and Serra-Garcia, M. (2013). Measuring lying aversion. *Journal of Economic Behavior and Organization*, 93:293–300.
- Gneezy, U., Saccardo, S., Serra-Garcia, M., and van Veldhuizen, R. (2020). Bribing the self. *Games and Economic Behavior*, 120:311–324.
- Godager, G. and Wiesen, D. (2013). Profit or patients' health benefit? Exploring the heterogeneity in physician altruism. *Journal of Health Economics*, 32(6):1105–1116.
- Green, E. P. (2014). Payment systems in the healthcare industry: An experimental study of physician incentives. *Journal of Economic Behavior & Organization*, 106:367–378.
- Grossman, Z. (2014). Strategic Ignorance and the Robustness of Social Preferences. *Management Science*, 60(11):2659–2665.
- Grossman, Z. and van der Weele, J. J. (2017). Self-image and willful ignorance in social decisions. *Journal of the European Economic Association*, 15(1):173–217.

- Hennig-Schmidt, H., Selten, R., and Wiesen, D. (2011). How payment systems affect physicians' provision behaviour-An experimental investigation. *Journal of Health Economics*, 30(4):637–646.
- Hennig-Schmidt, H. and Wiesen, D. (2014). Other-regarding behavior and motivation in health care provision: An experiment with medical and non-medical students. *Social Science & Medicine*, 108:156–165.
- Inderst, R. and Ottaviani, M. (2009). Misselling through agents. *American Economic Review*, 99(3):883–908.
- Inderst, R. and Ottaviani, M. (2012a). Competition through Commissions and Kick-backs. *American Economic Review*, 2012(2):780–809.
- Inderst, R. and Ottaviani, M. (2012b). Financial advice. *Journal of Economic Literature*, 50(2):494–512.
- Kajackaite, A. (2015). If I close my eyes, nobody will get hurt: The effect of ignorance on performance in a real-effort experiment. *Journal of Economic Behavior & Organization*, 116:518–524.
- Kajackaite, A. and Gneezy, U. (2017). Incentives and cheating. *Games and Economic Behavior*, 102:433–444.
- Kartik, N. (2009). Strategic communication with lying costs. *The Review of Economic Studies*, 76:1359–1395.
- Kerschbamer, R., Neururer, D., and Sutter, M. (2016). Insurance coverage of customers induces dishonesty of sellers in markets for credence goods. *Proceedings of the National Academy of Sciences*, 113(27):7454–7458.
- Kerschbamer, R., Neururer, D., and Sutter, M. (2019). Credence Goods Markets and the Informational Value of New Media: A Natural Field Experiment. Technical report, University of Innsbruck.
- Kerschbamer, R., Sutter, M., and Dulleck, U. (2017). How Social Preferences Shape Incentives in (Experimental) Markets for Credence Goods. *The Economic Journal*, 127(600):393–416.
- Khalmetski, K. and Sliwka, D. (2019). Disguising lies - Image Concerns and Partial Lying in Cheating Games. *American Economic Journal: Microeconomics*, 11(4):79–110.
- Kirchler, M., Huber, J., Stefan, M., and Sutter, M. (2016). Market design and moral behavior. *Management Science*, 62(9):2615–2625.

- Konow, J. (2000). Fair Shares: Accountability and Cognitive Dissonance in Allocation Decisions. *American Economic Review*, 90(4):1072–1091.
- Lerner, J. S. and Tetlock, P. E. (1999). Accounting for the effects of accountability.
- Liu, T. (2011). Credence Goods Markets With Conscientious And Selfish Experts. *International Economic Review*, 52(1):227–244.
- Mimra, W., Rasch, A., and Waibel, C. (2016). Price competition and reputation in credence goods markets: Experimental evidence. *Games and Economic Behavior*, 100:337–352.
- Momsen, K. (2021). Recommendations in credence goods markets with horizontal product differentiation. *Journal of Economic Behavior & Organization*, 183:19–38.
- Momsen, K. and Ohndorf, M. (2020a). Expressive voting vs. self-serving ignorance. Technical report, University of Innsbruck.
- Momsen, K. and Ohndorf, M. (2020b). When do people exploit moral wiggle room? an experimental analysis of information avoidance in a market setup. *Ecological Economics*, 169:106479.
- Momsen, K. and Ohndorf, M. (2022). Information avoidance, selective exposure, and fake (?) news: Theory and experimental evidence on green consumption. *Journal of Economic Psychology*, 88:102457.
- Nelson, P. (1970). Information and Consumer Behavior. *Journal of Political Economy*, 78(2):311–329.
- Onwezen, M. C., Antonides, G., and Bartels, J. (2013). The norm activation model: An exploration of the functions of anticipated pride and guilt in pro-environmental behaviour. *Journal of Economic Psychology*, 39:141–153.
- Petras, K., Ten Oever, S., and Jansma, B. M. (2016). The effect of distance on moral engagement: Event related potentials and alpha power are sensitive to perspective in a virtual shooting task. *Frontiers in Psychology*, 6:2008.
- Pigors, M. and Rockenbach, B. (2016). Consumer Social Responsibility. *Management Science*, 62(11):3123–3137.
- Pitchik, C. and Schotter, A. (1987). Honesty in a Model of Strategic Information Transmission. *American Economic Review*, 77(5):1032–1036.
- Rabin, M. (1994). Cognitive dissonance and social change. *Journal of Economic Behavior and Organization*, 23(2):177–194.

- Sandel, M. J. (2012). *What Money Can't Buy: The Moral Limits of Markets*. New York: Farrar, Straus and Giroux.
- Sassenberg, K., Boos, M., and Rabung, S. (2005). Attitude change in face-to-face and computer-mediated communication: Private self-awareness as mediator and moderator. *European Journal of Social Psychology*, 35(3):361–374.
- Schneider, T., Meub, L., and Bizer, K. (2021). Consumer information in a market for expert services: Experimental evidence. *Journal of Behavioral and Experimental Economics*, 94:101754.
- Scholl, A., Sassenberg, K., Zapf, B., and Pummerer, L. (2020). Out of sight, out of mind: Power-holders feel responsible when anticipating face-to-face, but not digital contact with others. *Computers in Human Behavior*, 112:106472.
- Schwartz, S. H. (1977). Normative influences on altruism. In *Advances in Experimental Social Psychology*, volume 10, pages 221–279. Elsevier.
- Sohn, S., Chung, H. C., and Park, N. (2019). Private Self-Awareness and Aggression in Computer-Mediated Communication: Abusive User Comments on Online News Articles. *International Journal of Human-Computer Interaction*, 35(13):1160–1169.
- Sülzle, K. and Wambach, A. (2005). Insurance in a market for credence goods. *Journal of Risk and Insurance*, 72(1):159–176.
- Sutter, M., Huber, J., Kirchler, M., Stefan, M., and Walzl, M. (2020). Where to look for the morals in markets? *Experimental Economics*, 23(1):30–52.
- Wolinsky, A. (1993). Competition in a Market for Informed Experts' Services. *The RAND Journal of Economics*, 24(3):380–398.
- Wolinsky, A. (1995). Competition in Markets for Credence Goods. *Journal of Institutional and Theoretical Economics (JITE)*, 151(1):117–131.

Appendix

A Instructions

Instructions

In this study, participants take the role of buyers and sellers of virtual products. Sellers offer two different types of virtual products, A and B.

The payoffs are displayed in Cents.

The seller's payoff depends on the product that they can sell. The seller earns a fixed payoff of 100 Cents plus a variable bonus that ranges from 90 to 130 Cents and corresponds to the payment associated with the market transaction.

In total, the seller makes 10 decisions of which one will be randomly chosen to determine the bonus payment. The bonus payment will be paid out as soon as a buyer is matched to the seller's decision in a separate session.

The buyer's payoff depends on the product that they purchase as well as the state of the world. There are two possible states of the world, both are equally likely. **The buyer does not know the state of the world.** The seller, in contrast, can observe the state of the world and the resulting payoffs for the buyer.

The seller decides which of the two products to sell to the buyer. **The buyer does not have the option to choose the other product.** The buyer earns a fixed payoff of 100 Cents plus a variable payoff that ranges from 20 to 120 Cents and depends on the purchased product and the state of the world.

Next

Figure 5: Screenshot - Instructions, sell

Instructions

In this study, participants take the role of buyers and sellers of virtual products. Sellers offer two different types of virtual products, A and B.

The payoffs are displayed in Cents.

The seller's payoff depends on the product that they can sell. The seller earns a fixed payoff of 100 Cents plus a variable bonus that ranges from 90 to 130 Cents and corresponds to the payment associated with the market transaction.

In total, the seller makes 10 decisions of which one will be randomly chosen to determine the bonus payment. The bonus payment will be paid out as soon as a buyer is matched to the seller's decision in a separate session.

The buyer's payoff depends on the product that they purchase as well as the state of the world. There are two possible states of the world, both are equally likely. **The buyer does not know the state of the world.** The seller, in contrast, can observe the state of the world and the resulting payoffs for the buyer.

The seller decides which of the two products to recommend to the buyer. **The buyer can follow the seller's recommendation, but they can also choose differently.** The buyer earns a fixed payoff of 100 Cents plus a variable payoff that ranges from 20 to 120 Cents and depends on the purchased product and the state of the world.

In each of the 10 decisions, the seller will be matched with a different buyer.

Next

Figure 6: Screenshot - Instructions, recommend

B Decision screens

Market transaction

You are a seller. You offer two different products A and B.

Your client's payoffs from the products depend on the state of the world; we are in state H. **Note that your client cannot observe the state of the world.** The purchased product and the state of the world determine their payoff.

Your client does not have the option to choose the other product. Which product do you sell to your client?

	Your payoff	Your client's payoff in state L	Your client's payoff in state H (current state)	
Product A	90	20	120	<input type="button" value="Sell A"/>
Product B	130	100	100	<input type="button" value="Sell B"/>

Figure 7: Screenshot - Full information, sell

Market transaction

You are a seller. You offer two different products A and B.

Your client's payoff from the products depend on one of two possible states of the world ('H' or 'L'); both states are equally likely. **Note that your client cannot observe the state of the world.** The purchased product and the state of the world determine their payoff.

Revealing the state of the world is **optional**: You can reveal the state of the world by clicking the button below, but you can also make your selling decision without knowing the state of the world.

Your client does not have the option to choose the other product. Which product do you sell to your client?

	Your payoff	Your client's payoff in state L (50%)	Your client's payoff in state H (50%)	
Product A	90	20	120	<input type="button" value="Sell A"/>
Product B	130	100	100	<input type="button" value="Sell B"/>
		<input type="button" value="Optional: Reveal the state of the world"/>		

Figure 8: Screenshot - Hidden information, sell

Market transaction

You are a seller. You offer two different products A and B.

Your client's payoff from the products depend on the state of the world; we are in state H. **Note that your client cannot observe the state of the world.** The purchased product and the state of the world determine their payoff.

For this situation, you will be randomly matched with one out of 50 clients.

Your client can follow, but does not have to follow your recommendation. Which product do you recommend to your client?

	Your payoff	Your client's payoff in state L	Your client's payoff in state H (current state)	
Product A	100	20	100	Recommend A
Product B	120	100	20	Recommend B

Figure 9: Screenshot - Full information, recommend

Market transaction

You are a seller. You offer two different products A and B.

Your client's payoff from the products depend on one of two possible states of the world ('H' or 'L'); both states are equally likely. **Note that your client cannot observe the state of the world.** The purchased product and the state of the world determine their payoff.

Revealing the state of the world is **optional**: You can reveal the state of the world by clicking the button below, but you can also make your recommendation decision without knowing the state of the world.

For this situation, you will be randomly matched with one out of 50 clients.

Your client can follow, but does not have to follow your recommendation. Which product do you recommend to your client?

	Your payoff	Your client's payoff in state L (50%)	Your client's payoff in state H (50%)	
Product A	130	20	120	Recommend A
Product B	90	100	100	Recommend B
Optional: Reveal the state of the world				

Figure 10: Screenshot - Hidden information, recommend

University of Innsbruck - Working Papers in Economics and Statistics
Recent Papers can be accessed on the following webpage:

<https://www.uibk.ac.at/eeecon/wopec/>

- 2022-10 **Katharina Momsen, Markus Ohndorf:** Seller Opportunism in Credence Good Markets – The Role of Market Conditions
- 2022-09 **Christoph Huber, Michael Kirchler:** Experiments in Finance – A Survey of Historical Trends
- 2022-08 **Tri Vi Dang, Xiaoxi Liu, Florian Morath:** Taxation, Information Acquisition, and Trade in Decentralized Markets: Theory and Test
- 2022-07 **Christoph Huber, Christian König-Kersting:** Experimenting with Financial Professionals
- 2022-06 **Martin Gächter, Martin Geiger, Elias Hasler:** On the structural determinants of growth-at-risk
- 2022-05 **Katharina Momsen, Sebastian O. Schneider:** Motivated Reasoning, Information Avoidance, and Default Bias
- 2022-04 **Silvia Angerer, Daniela Glätzle-Rützler, Philipp Lergetporer, Thomas Rittmannsberger:** How does the vaccine approval procedure affect COVID-19 vaccination intentions?
- 2022-03 **Robert Böhm, Cornelia Betsch, Yana Litovsky, Philipp Sprengholz, Noel Brewer, Gretchen Chapman, Julie Leask, George Loewenstein, Martha Scherzer, Cass R. Sunstein, Michael Kirchler:** Crowdsourcing interventions to promote uptake of COVID-19 booster vaccines
- 2022-02 **Matthias Stefan, Martin Holmén, Felix Holzmeister, Michael Kirchler, Erik Wengström:** You can't always get what you want—An experiment on finance professionals' decisions for others
- 2022-01 **Toman Barsbai, Andreas Steinmayr, Christoph Winter:** Immigrating into a Recession: Evidence from Family Migrants to the U.S.
- 2021-32 **Fanny Dellinger:** Housing Support Policies and Refugees' Labor Market Integration in Austria
- 2021-31 **Albert J. Menkveld, Anna Dreber, Felix Holzmeister, Jürgen Huber, Magnus Johannesson, Michael Kirchler, Sebastian Neusüss, Michael Razen, Utz Weitzel and et al:** Non-Standard Errors

- 2021-30 **Toman Barsbai, Victoria Licuanan, Andreas Steinmayr, Erwin Tiongson, Dean Yang:** Information and Immigrant Settlement
- 2021-29 **Natalie Struwe, Esther Blanco, James M. Walker:** Competition Among Public Good Providers for Donor Rewards
- 2021-28 **Stjepan Srhoj, Melko Dragojević:** Public procurement and supplier job creation: Insights from auctions
- 2021-27 **Rudolf Kerschbamer, Regine Oexl:** The effect of random shocks on reciprocal behavior in dynamic principal-agent settings
- 2021-26 **Glenn E. Dutcher, Regine Oexl, Dmitry Ryvkin, Tim Salmon:** Competitive versus cooperative incentives in team production with heterogeneous agents
- 2021-25 **Anita Gantner, Regine Oexl:** Respecting Entitlements in Legislative Bargaining - A Matter of Preference or Necessity?
- 2021-24 **Silvia Angerer, E. Glenn Dutcher, Daniela Glätzle-Rützler, Philipp Lergetporer, Matthias Sutter:** The formation of risk preferences through small-scale events
- 2021-23 **Stjepan Srhoj, Dejan Kovač, Jacob N. Shapiro, Randall K. Filer:** The Impact of Delay: Evidence from Formal Out-of-Court Restructuring
- 2021-22 **Octavio Fernández-Amador, Joseph F. Francois, Doris A. Oberdabernig, Patrick Tomberger:** Energy footprints and the international trade network: A new dataset. Is the European Union doing it better?
- 2021-21 **Felix Holzmeister, Jürgen Huber, Michael Kirchler, Rene Schwaiger:** Nudging Debtors to Pay Their Debt: Two Randomized Controlled Trials
- 2021-20 **Daniel Müller, Elisabeth Gsottbauer:** Why Do People Demand Rent Control?
- 2021-19 **Alexandra Baier, Loukas Balafoutas, Tarek Jaber-Lopez:** Ostracism and Theft in Heterogeneous Groups
- 2021-18 **Zvonimir Bašić, Parampreet C. Bindra, Daniela Glätzle-Rützler, Angelo Romano, Matthias Sutter, Claudia Zoller:** The roots of cooperation
- 2021-17 **Silvia Angerer, Jana Bolvashenkova, Daniela Glätzle-Rützler, Philipp Lergetporer, Matthias Sutter:** Children's patience and school-track choices several years later: Linking experimental and field data
- 2021-16 **Daniel Gründler, Eric Mayer, Johann Scharler:** Monetary Policy Announcements, Information Shocks, and Exchange Rate Dynamics
- 2021-15 **Sebastian Bachler, Felix Holzmeister, Michael Razen, Matthias Stefan:** The Impact of Presentation Format and Choice Architecture on Portfolio Allocations: Experimental Evidence

- 2021-14 **Jeppe Christoffersen, Felix Holzmeister, Thomas Plenborg:** What is Risk to Managers?
- 2021-13 **Silvia Angerer, Daniela Glätzle-Rützler, Christian Waibel:** Trust in health care credence goods: Experimental evidence on framing and subject pool effects
- 2021-12 **Rene Schwaiger, Laura Hueber:** Do MTurkers Exhibit Myopic Loss Aversion?
- 2021-11 **Felix Holzmeister, Christoph Huber, Stefan Palan:** A Critical Perspective on the Conceptualization of Risk in Behavioral and Experimental Finance
- 2021-10 **Michael Razen, Alexander Kupfer:** Can increased tax transparency curb corporate tax avoidance?
- 2021-09 **Changxia Ke, Florian Morath, Anthony Newell, Lionel Page:** Too big to prevail: The paradox of power in coalition formation
- 2021-08 **Marco Haan, Pim Heijnen, Martin Obradovits:** Competition with List Prices
- 2021-07 **Martin Dufwenberg, Olof Johansson-Stenman, Michael Kirchler, Florian Lindner, Rene Schwaiger:** Mean Markets or Kind Commerce?
- 2021-06 **Christoph Huber, Jürgen Huber, and Michael Kirchler:** Volatility Shocks and Investment Behavior
- 2021-05 **Max Breitenlechner, Georgios Georgiadis, Ben Schumann:** What goes around comes around: How large are spillbacks from US monetary policy?
- 2021-04 **Utz Weitzel, Michael Kirchler:** The Banker's Oath And Financial Advice
- 2021-03 **Martin Holmen, Felix Holzmeister, Michael Kirchler, Matthias Stefan, Erik Wengström:** Economic Preferences and Personality Traits Among Finance Professionals and the General Population
- 2021-02 **Christian König-Kersting:** On the Robustness of Social Norm Elicitation
- 2021-01 **Laura Hueber, Rene Schwaiger:** Debiasing Through Experience Sampling: The Case of Myopic Loss Aversion.

University of Innsbruck

Working Papers in Economics and Statistics

2022-10

Katharina Momsen, Markus Ohndorf

Seller Opportunism in Credence Good Markets – The Role of Market Conditions

Abstract

We report the results of an experiment to systematically investigate the influence of different settings in credence good markets on opportunism in the sellers' decisions. We find that, as predicted by a cognitive dissonance model, the specific choice of the design features might be less innocuous than generally presumed: sellers' decisions made under a direct sales regime are significantly more opportunistic than purchase recommendations. Furthermore, average choices are more opportunistic when a costless diagnosis is required to assess the buyer's needs — sellers exploit moral wiggle room by avoiding information. Yet, this effect is only present for purchase recommendations, not direct sales. Both of these effects significantly affect market efficiency. Generally, the parametrization of the decision problem has a strong influence on opportunism, as predicted. Here, we find that sellers tend to overtreat and buyers self-select into overtreatment.

ISSN 1993-4378 (Print)

ISSN 1993-6885 (Online)