

Revealing Good Deeds: Disclosure of Social Responsibility in Competitive Markets

Sören Harrs*

Bettina Rockenbach[†]

Lukas M. Wenner[‡]

February 2022

Abstract

We experimentally study competitive markets with socially responsible production. Our main focus is on the producers' decision whether or not to reveal the degree of social responsibility of their product. Compared to two benchmark cases where either full transparency is enforced or no disclosure is possible, we show that voluntary and costless disclosure comes close to the full transparency benchmark. However, when the informational content of disclosure is imperfect, social responsibility in the market is significantly lower than under full transparency. Our results highlight an important role for transparent and standardized information about social externalities.

Keywords: Social responsibility; market experiments; disclosure; consumer behavior

JEL Classification Numbers: C92; D90; D62; M14

*Department of Economics, University of Cologne, Albertus Magnus Platz, 50923 Cologne, e-mail: harrs@wiso.uni-koeln.de.

[†]Department of Economics, University of Cologne, Albertus Magnus Platz, 50923 Cologne, e-mail: bettina.rockenbach@uni-koeln.de.

[‡]Department of Economics, University of Cologne, Albertus Magnus Platz, 50923 Cologne, e-mail: lukas.wenner@uni-koeln.de.

The replication material for the study is available at <https://doi.org/10.17605/OSF.IO/KZ2FH>. Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2126/1– 390838866 and by the Center for Social and Economic Behavior (C-SEB). We thank Jana Friedrichsen, Dorothea Kübler, Adriaan Soetevent and conference participants at ESA Europe 2017 (Vienna) for helpful comments. Viet A Nguyen and Hanan Iqbal provided valuable help in running the experiments.

1 Introduction

The social responsibility (SR) of products and their production processes is a central topic of many public controversies. Prime examples are the accusation of deficiencies in working conditions (e.g., in garment factories), damaging environmental effects (e.g., in CO₂ emissions), or a lack of sustainability in production (e.g., in using non-recyclable plastics). While there seems to be an increasing awareness of consumers about the wider social impact of the purchase of many products, suggestive evidence points at overall levels of SR in markets being fairly low.¹

The existing literature, however, shows that it is not that consumers do not care at all about SR, but instead exhibit a significant increase in their willingness-to-pay for these types of products.² Then how does this square with the idea that in many industries SR products often only have very small market shares? If consumers care about the social responsibility of the purchased products, this should be a relevant dimension for firms to compete in, and firms offering SR products should have a competitive advantage. This argument, however, relies on the assumption that consumers can assess the social impact of a product. Yet, typically SR is not a directly visible characteristic of the product, but one that needs to be actively revealed by the seller, like, for example, the working conditions in the production process. Although in some domains revelation of the SR in production through producer declarations seems an appropriate solution (e.g., declaration of a piece of furniture to be free of tropical timber), in many cases incentives for providing non-truthful specifications are strong and hardly verifiable (e.g., due to complex opaque production chains). In such cases, labels certifying SR in production may be a powerful solution. The multiplicity of labels has tremendously increased in recent years. Some labels address single issues, like the energy efficiency class of an electric appliance, while others cover multiple facets. For example, fair trade as well as animal welfare are two dimensions of social responsibility where a large number of different labels for seemingly similar issues exist.³ Arguably, if the externalities of products are measured in various different ways by considering various different facets, this

¹Taking the market for organic foods as one example, in the United States, the share of organic foods compared to total food sales in 2017 is estimated to be at 5.5% (OTA, 2018). In Europe, Denmark has the highest share of organic retails sales at 13.3% of the total market (FiBL, 2017).

²See, for example, Bartling et al. (2015; 2019; 2020b) or Pigors and Rockenbach (2016) for laboratory evidence and Hainmueller et al. (2015) or Elfenbein and McManus (2010) for field evidence.

³For instance, for the United States, the magazine *Consumer Reports* presents a detailed guide navigating through the different fair trade labels for coffee, identifying five trustworthy ones and, in addition, warning consumers about a number of meaningless claims. (Dabney, 2017). The German Federation for the Environment and Nature Conservation (BUND) lists seventeen different labels that are currently used to certify ecological and sustainable standards in production (BUND, 2019).

considerably lowers the informational content of those labels for consumers.

In this paper, we experimentally study the emergence of SR in competitive markets where SR is not a directly visible characteristic of the product, but information about the SR of the product can be voluntarily provided by its seller. In the experiment, our proxy for SR is a monetary transfer to the charity *unicef*. Each seller chooses the price at which he offers the product, and an amount (also referred to as “donation” below) to be transferred to *unicef* in case he sells the product to a consumer. Donations do not influence the consumer’s monetary benefit from the product, but the vast majority of participants in our experiment consider *unicef* a charity worth supporting. Higher donations incur higher monetary costs for the seller.⁴

To get a nuanced picture we study various scenarios in our experimental treatments. Our benchmark cases are that customers are either completely ignorant (NO INFO) or fully informed (FULL INFO) about the donation associated with the product (i.e. the SR of the product). The degree of information (either none or full) is externally enforced and neither producers nor customers can change this. The FULL INFO case mirrors a situation in which, for example, regulation requires an unambiguous declaration of the SR characteristic of the product, e.g. the specification of the energy efficiency class of a light bulb or a fridge.

Three further CHOICE treatments mimic situations in which producers can choose to reveal the SR information of their product to customers (e.g., by certification through a label), but revelation is potentially imperfect. The CHOICE treatments vary the degree of the imperfection of this revelation. In treatment CHOICE-100 customers unambiguously see the information whenever a producer decides to reveal the SR information. This resembles the case of a (standardized) label and customers that are aware of what this label certifies (e.g., declaration of the absence of tropical timber). In treatments CHOICE-85 and CHOICE-60 the producer also decides whether or not to reveal the SR information, but - in case of a revelation decision - both the revelation decision and the revealed SR information may not reach the consumer (in CHOICE-85 there is a 85 percent chance that the revelation decision and SR information reaches consumers and there is a 60 percent chance in CHOICE-60). These two treatments are designed to map situations in which the producer faces an uncertainty whether his SR information reaches the consumers. Reasons for this uncertainty may be that consumers do not get the informational content of the provided information, e.g., due to a lack of standardization or limited attention of consumers. Hence, in treatments CHOICE-85 and CHOICE-60, a consumer facing non-disclosed information does not know

⁴Using transfers to *unicef* as a proxy for SR has been done previously, e.g., in Sutter et al. (2020) and Kirchler et al. (2016). We justify our choice of using this proxy for SR further in Section 3 when presenting the details of our experimental design

whether the producer chose not to disclose or whether he disclosed, but the information was not transmitted. This means that while strategically sophisticated buyers should interpret undisclosed donations in CHOICE-100 as zero donations, this is no longer the case in CHOICE-85 and CHOICE-60. The unreliability of disclosure provides strategic incentives to producers to offer products with a lower level of social responsibility compared to CHOICE-100 where buyers can rely on non-disclosure being fully reliable.

In a nutshell, our results show that disclosure of SR information has large and positive effects compared to a regime where producers are unable to disclose the social responsibility of their products to consumers (NO INFO). We find that when voluntary disclosure provides unambiguous information for customers (CHOICE-100), the level of social responsibility comes close to the outcome under full and externally enforced disclosure (FULL INFO). While this is, in general, positive news when trying to establish more socially responsible markets, our results for the cases where disclosure is imperfect show the limitations of such an approach. When producers cannot be sure that the provided SR information has full informational content for consumers (CHOICE-85 and CHOICE-60), SR in production is significantly lower compared to FULL INFO. Yet, since our data does not provide an unambiguous statistical difference between the three choice treatments, the sensitivity of social responsibility to changes in the disclosure probability cannot be precisely determined. Furthermore, we show that producers benefit from more opaque information environments via higher profits. They produce products at lower cost, while selling them with a higher mark-up, taking advantage of the reduced competition in the market.

Our results have important implications for the understanding of markets with a social externality. In particular, they can provide guidance for policy makers and regulators interested in increasing socially responsible production or products. This is true in particular for environments where there is scope for “demand-driven” social responsibility, i.e., markets where consumers care about the social impact of the product they purchase and have enough influence to affect sellers’ production choices. Then, competitive markets can deliver high levels of social responsibility, but only when transparency and clarity of the latter is enforced. Since such strict policies may be impractical to implement, we show that voluntary disclosure, which requires little regulatory intervention, seems to be a viable alternative to also achieve high levels of SR. However, our results also indicate that this is not true if producers can disclose the social responsibility of their products only in a way that is unreliable. This provides a rationale for coherent labels with standardized criteria. In many relevant industries, however, this is not the case, as the examples provided above show. Our results are consistent with markets for such products displaying low production standards for a majority of the products sold. They also provide a clear reason why neither forced

disclosure nor more informative and standardized labels are in the interest of the producers: Producers benefit from the reduced transparency through higher profits, and thus have strong incentives to actively oppose such labels.

2 Related Literature

Our paper adds to an increasing body of literature in economics which studies the determinants of moral behavior, see e.g., Irlenbusch and Villeval (2015) for an overview. More specifically, we focus on the emergence of social responsibility in markets. To our knowledge, there is no paper within the experimental literature which explicitly analyzes, as our paper does, the firms' decision whether or not to reveal the level of social responsibility of the offered product.

Bartling et al. (2015) study a laboratory market where products with and without an externality, represented by the earnings of a third party, can be traded. They find that behavior in those markets reflects a concern for social responsibility, demonstrated through market shares for the socially responsible product of up to 50%. Using the same basic paradigm, Bartling et al. (2019) show that these results do not depend on whether social harm is concentrated on a single subject or diffused among many, while Bartling et al. (2020b) show that an exogenous increase in consumer incomes increases social responsibility in the market. Using, like us, a charity rather than another lab participant representing the externality, Rode et al. (2008) also provide evidence that consumers are willing to pay a price premium for socially responsible products, thus avoiding them being crowded out. These findings are in accordance with our results in treatment FULL INFO: when consumers are fully aware of the social impact of their purchase decisions, market outcomes reflect their concerns for socially responsible behavior.

A closely related set of papers contains a comparison across market conditions which is similar to the relation between the treatments NO INFO and FULL INFO in this paper. Pigors and Rockenbach (2016), in a setting similar to Bartling et al. (2015), also show that once the social dimension becomes known to buyers rather than being private information of the sellers, social responsibility in the market increases. In Feicht et al. (2016), when they face neither commitment nor a possibility that deviating behavior is observed, sellers generate lower donations compared to the case where there is full commitment to donations as promised to buyers at the point of purchase.

Etilé and Teyssier (2016) study the roles of labels in an experimental market. They compare treatments where sellers can acquire labels signifying a minimum amount of donations

to a charity to a treatment where no labels can be acquired. They find that only if such labels are credible, i.e., sellers cannot lie, the use of such labels increases social responsibility. Our paper differs from theirs in, at least, two important aspects. First, they cannot say anything about the effectiveness of labels compared to an institutional regime which forces disclosure of the amount of social responsibility. Neither do they allow for sellers to convey information to buyers beyond whether or not a minimum standard is met. Moreover, our design introduces the feature that labels are often complicated and difficult to transmit to consumers, something which is both realistic and relevant for policy makers in so far as markets outside the lab are concerned.⁵

On a broader level, we add to a literature which studies under which conditions market forces are beneficial in generating socially and morally desirable outcomes. Falk and Szech (2013) argue that the presence of markets reduces moral behavior, as indicated by lower prices in markets than in individual decisions for the same externalities (saving the life of a mouse). Yet in a recent paper Bartling et al. (2020a) show that neither Falk and Szech’s data nor their own data support the claim that market interaction erodes moral behavior, but instead it is repeated play that causes the erosion of moral values. Kirchler et al. (2016) also compare individual decision making with a double-auction market and identify the threat of punishment as an effective intervention to increase moral behavior. Sutter et al. (2020) provide evidence that markets with an externality lead to fewer trades, while the effect on prices is ambiguous. The latter two papers use, like us, donations to *unicef* as a proxy for moral behavior. Pigors and Rockenbach (2016) compare monopoly to duopoly markets and show that competition has positive effects on the emergence of more moral behavior. In light of this literature, our design can be interpreted as varying the intensity of competition along the “informativeness-dimension” of the social externality. If there is competition both on the price and the charity dimension, social responsibility increases compared to cases where the latter dimension is not, or only partly, relevant for competitors.

Beyond the literature on social responsibility, our paper also speaks to a literature on information disclosure in markets. Recent evidence (e.g., Jin et al., 2021; Benndorf et al., 2015; Deversi et al., 2021) shows that people have difficulties interpreting non-disclosure as sufficiently negative information, severely limiting the kind of unraveling predicted when agents were sophisticated rather than boundedly rational. Given that in our CHOICE treatments buyers face a similar task, namely to infer the donation associated with undisclosed

⁵The paper by Engelmann et al. (2018) addresses the argument of external validity of experimental studies of socially responsible behavior in markets. Reassuringly, they demonstrate that buyers who show high levels of social responsibility in experimental markets, also have higher willingness-to-pay for real fair trade chocolate over conventional one.

offers, it is worth pointing out that if disclosure is sufficiently unambiguous, our results show that buyers behave as if they can infer donations relatively well. An important difference of our paper, also compared to the theoretical literature on disclosure in the presence of boundedly rational consumers (e.g., Ispano and Schwardmann, 2018), is that we implement a setting where sellers can choose their donation level, i.e., their product’s quality. Hence, this may be an important driver of the diverging findings, in particular as in most markets quality is endogenous, at least in the medium- to long-run.⁶

3 Experimental Design

Participants in our experiment interact on markets with four sellers and two buyers. Each participant receives an endowment of $e = 20$ [points] in each period. We implement posted-offer markets where buyers arrive sequentially. That is, in each of the 30 periods, one buyer per market is selected at random and faces four offers, one from each seller. This buyer then chooses one (or none) of the four offers and then the second buyer chooses from the remaining offers. A buyer who purchases a product obtains a benefit of $v = 100$ and pays a price p to the seller. The seller chooses the price $p \in [0, 120]$ at which he offers the product, and an amount $d \in \{0, 10, 20, \dots, 100\}$ to be transferred to *unicef* in case he sells the product to a buyer. These donations do not influence the buyer’s monetary benefit from the product (v), yet higher donations incur higher monetary costs $c(d)$ for the seller, as shown by the following table:⁷

d	100	90	80	70	60	50	40	30	20	10	0
$c(d)$	29	27	25	23	21	18	15	12	8	4	0

Thus, monetary payoffs of buyers and sellers are:

$$\Pi^b = \begin{cases} e & \text{if buyer does not buy} \\ e + v - p & \text{if buyer buys} \end{cases} \quad (1)$$

⁶A further, though somewhat more distant connection can be drawn from our work to the literature on experimental “lemon markets” where quality choice is endogenous (e.g., Holt and Sherman, 1990; Henze et al., 2015). This is because if one interprets SR production as the quality of a product, our NO INFO and FULL INFO treatments can be understood as capturing similar ideas as in studying markets with full or asymmetric information on quality. Our finding that social responsibility is lower in NO INFO than in FULL INFO is consistent with the emergence of a “lemons outcome” found in this literature.

⁷This cost function exhibits decreasing average cost which can be seen as capturing some implicit fixed cost in socially responsible production. For the sake of experimental simplicity, we did not implement fixed costs explicitly.

$$\Pi^s = \begin{cases} e & \text{if buyer does not buy} \\ e + p - c(d) & \text{if buyer buys} \end{cases} \quad (2)$$

From a design perspective, implementing the social externality by means of a monetary transfer to *unicef* is attractive for at least two reasons. First, the donation to *unicef* is - as confirmed through a questionnaire at the end of the experiment (see Section 5.3) - clearly viewed by the parties trading on the market as a positive externality. Hence, it can suitably capture, in a stylized manner, the choice of a firm how socially responsible the offered product is. Second, money transferred to UNICEF is not a transfer between market participants, but something that affects people not involved in the transaction, just like many externalities outside of the laboratory, such as environmental effects. Moreover, we view the possibility that participants could, in principle, also donate to *unicef* outside of the experiment as a feature which increases the external validity of the experiment because in many real-world settings consumers and firms can choose between buying more socially responsible products on a market, like a low CO₂ emission vehicle, or can compensate higher emissions through other means.⁸

3.1 Treatments

We implement five treatments which vary the information about the donation. There are two benchmark treatments, one where information on each product’s associated donation is fully disclosed and one where the information is fully hidden: In the NO INFO treatment, buyers do not observe whether or not donations are made in association with the products offered. The only information buyers have is the price of each product. They also cannot observe this information at a later stage. Neither do sellers observe the donation decisions of the other sellers. In the FULL INFO treatment, buyers (and the other sellers) observe the donation decisions and levels of all products as well as the products’ prices.

In addition, we run three CHOICE treatments where information can be disclosed by sellers voluntarily. The three treatments differ in the informational content provided. In treatment CHOICE-100, whenever sellers decide to reveal the level of donation d , customers unambiguously see the production information, i.e., d is revealed with probability 100%. In the treatments CHOICE-85 and CHOICE-60, the revelation decision of the seller is imple-

⁸Due to the fact that a donation d has a cost of $c(d) < d$ for all d for a seller, it is however more efficient for market participants as a whole to generate donations in the experiment rather than outside of the experiment. This is not necessary for the external validity of the design, but makes it less attractive for participants to substitute generating donations in the experiment by donating afterwards outside of the experiment.

mented with probability 85% and 60%, respectively. Thus, when a seller in a given period decides to reveal his donation to the buyers, a random draw by the computer determines whether the decision to reveal and the donation associated with the product are visible to the buyers (and the other sellers) or not. In all three CHOICE treatments, the decision not to reveal is always implemented with certainty.⁹

3.2 Procedures

The markets in which buyers and sellers interact retain the same composition throughout, that is, there is no re-matching across periods, and participants are aware of this. After all decisions within a round have been made, all participants are informed about their earnings. In addition, throughout the experiment, every subject has access to the full information of his/her past behavior, i.e., prices, earnings and (if known) generated donations. Also, all buyers and sellers are identified through letters and behavior can thus be traced across rounds.¹⁰

To calculate the total earnings of the participants, the earnings from all periods are added up and converted into Euros at the rate of 130 points = 1 EUR. The same conversion rate applies to donations. These donations are transferred to *unicef* at the day of the experiment. Since donations are both the result of market interaction and sometimes private information (and we wanted to keep this private), we provide all participants with a session average of the donations. This strikes a balance between not revealing much about the individual behavior but at the same time maintaining credibility about our procedures. We furthermore send proof of the donation to all participants (unless they told us they preferred not to receive such an email).

We implemented 12 markets per treatment, leading to 360 participants overall. After the 30 periods of the main experiment, subjects answered a short questionnaire which elicited perceptions about *unicef* and donation behavior outside the lab (how often the person donated to charities in the past, and how much), as well as demographics such as age, gender and subject of study. The experiments were run at the Cologne Experimental Economics Laboratory (CLER) between May 2017 and October 2018. Recruiting was done via ORSEE

⁹As a robustness check we also ran a variant of treatment CHOICE-100 where we introduce a disclosure cost of 3 but at the same time shift the cost $c(d)$ down by 3 points. While the cost associated with disclosure in this treatment are identical to CHOICE-100, this is not the case for sellers who decide not to disclose. Hence, buyers could be more likely to believe that sellers do not disclose positive donations in order to save on the disclosure cost. In Section A.6 in the Appendix we show that behavior is similar to CHOICE-100 and discuss the implications of this result. We refrain from including it in the main text because due to the different disclosure costs, this treatment is not directly comparable to the rest.

¹⁰For a (translated) version of the experimental instructions, see Appendix B.

(Greiner, 2015) and the experiment was programmed in zTree (Fischbacher, 2007). Average earnings amounted to 11.64 EUR plus a 4 EUR show-up fee, and average contributions to *unicef* were 68.15 EUR per session.

In this experiment we are powered to detect effects of around one standard deviation (or larger) with 12 markets per treatment condition (MWU-tests, $\alpha = 0.05$, power=80%, see Appendix A.5 for details). To have sufficient power to be able to judge the (non-)difference between Choice-100 and Full Info (see Result 2), we conducted a replication study in March 2021 consisting of 27 markets each for the treatments FULL INFO and CHOICE-100. For a detailed description, including those aspects in which the replication study differs from the initial study, see Appendix A.5.

4 Predictions

In this section, we derive the predictions for our experiment. We start by considering treatments NO INFO and FULL INFO. These two treatments differ only in whether or not buyers can observe the donations chosen by sellers. In treatment NO INFO, buyers only see the prices sellers ask for. Buyers may believe that some of the higher prices may come from sellers with positive donations. Thus, buyers with social preferences may give up monetary payoff and accept higher prices. Yet, buyers are aware that higher prices are not a reliable signal of positive donations, because they might be an attempt of sellers to achieve higher profits. In contrast, sellers in treatment FULL INFO compete both on price and the level of social responsibility. A product will be the more attractive to a buyer the lower its price and - for a buyer with some form of social preferences that make her also value money transferred to the charity - the higher the donations associated with a purchase. That donations are always revealed to buyers should lead to a higher level of social responsibility in the market but also to higher prices (because sellers will need to cover the higher production costs) than in NO INFO.

Prediction 1 (NO INFO vs. FULL INFO). *Forcing disclosure of the donation leads to a higher level of social responsibility compared to the case where donations are not visible for buyers. Due to the increased production cost, market prices will be higher when donations are disclosed.*

Next, we consider the three CHOICE treatments. We start by focusing on the CHOICE-100 treatment, and then move on to CHOICE-85 and CHOICE-60. The crucial part of the analysis concerns the beliefs that buyers form when faced with offers that do not have the

donation level disclosed. Even though we do not specify a formal model to derive predictions, in the following we sketch a simple yet powerful unraveling intuition which we believe is sufficiently general to apply to our specific setting as well.

Specifically, we posit that a rational buyer holds the belief that the donation level associated with an undisclosed offer equals the average donation level set by all those sellers who do not disclose their donations. Now suppose that among the sellers who choose not to disclose their donation, a share $s < 1$ chooses a donation level $d' > 0$, while the rest chooses $d = 0$. Then, the belief a rational buyer holds about the donation level in undisclosed offers would be $sd' < d'$. But then, a seller who chooses d' is better off disclosing her donation level to separate from those sellers with $d = 0$ and thereby increasing the buyers' willingness to pay for her product. At the same time, if $s = 1$, i.e. all sellers who choose not to disclose set $d' > 0$ (and buyers believe this to be the case), there is a profitable deviation for a seller to reduce his donation level to $d = 0$ because this increases his profits due to the lower production cost without affecting the purchase decision of a buyer (since the donation decision is unobservable). This implies that there should not be any sellers who shroud offers with positive donations.

Importantly, if buyers do not hold fully rational beliefs, as shown in recent experimental work on disclosure games (e.g., Jin et al., 2021; Benndorf et al., 2015; Deversi et al., 2021), it might be possible that sellers could exploit such biased beliefs by profitably choosing more zero donations compared to FULL INFO, but not disclosing them to buyers. However, our setting differs from those disclosure games in a number of ways which means that we consider it an open question whether we should expect beliefs to be biased in our case.¹¹ We therefore state our prediction based on the assumption that buyers form rational beliefs.

Prediction 2 (CHOICE-100 vs. FULL INFO). *Making disclosure of donations voluntary leads to the same market outcome as when sellers are forced to reveal this information. Donations and market prices are identical in treatments CHOICE-100 and FULL INFO.*

Note, however, that this prediction does not tell us anything about the disclosure rates, i.e., the share of undisclosed offers in treatment CHOICE-100. Based on the intuition developed above, donations of zero which are disclosed and undisclosed donations should be treated by buyers in a similar manner, which in turn means that for sellers who choose a donation of zero it is irrelevant whether they disclose or not. What we can thus only say is that we do not expect the share of undisclosed offers in treatment CHOICE-100 to be larger than the share of offers with zero donations in treatment FULL INFO.

¹¹The most important difference is the fact that we consider a setting in which the informed party (sellers), selects the relevant quantity (donation) before deciding whether to disclose it or not. In the disclosure games cited above, the disclosure choice typically concerns a quantity which has been determined randomly.

When analyzing the cases where a seller’s decision to reveal the social component of the offer is not implemented with certainty, it is no longer the case that buyers can be sure that an undisclosed donation is due to sellers actively wanting to hide their donation. It could be the case that the seller chooses a high donation but the chance move prevented this donation from being visible to buyers. This means that it is a lot less straightforward for buyers to interpret undisclosed donations in treatments CHOICE-85 and CHOICE-60 compared to treatment CHOICE-100.

Intuitively, we thus predict that such an unreliability in disclosure allows sellers to strategically exploit this more opaque environment. As buyers have to entertain the possibility that an undisclosed donation may in fact also be a high donation, they should, compared to CHOICE-100 be more inclined to buy products with hidden donations. This, however, should make it a profitable strategy for sellers to (sometimes) choose a low donation and to not disclose. As a countervailing force to this mechanism, however, it also does not seem a profitable strategy for sellers to always choose to not disclose. Since the probability of disclosure is fairly high in both treatments, a seller who decides to disclose a high donation - and this decision actually reaches the buyers - is likely to have a competitive advantage, especially if competing offers have their donation hidden.

Taking these effects together, we predict that the possibilities for sellers to exploit the limited informational value of non-disclosure lead to lower donations compared to the FULL INFO benchmark. Furthermore, since, as outlined in the previous paragraph, sellers can use the option to disclose strategically in order to exploit the inherent unreliability, we expect that this softens competition between sellers and allows them to charge higher prices, compared to treatment FULL INFO. We do not make predictions as to whether these effects are larger in treatment CHOICE-85 or in CHOICE-60 as this would require a formal model which is beyond the scope of this paper.

Prediction 3 (CHOICE-85/CHOICE-60 vs. FULL INFO/CHOICE-100). *When donations are voluntary but donation decisions of sellers have limited informational value, donations are lower and prices are higher compared to when there is perfect information.*

5 Results

We divide the presentation of the results in several parts. First, in Section 5.1, we analyze how the treatments compare with respect to donations generated and (in the CHOICE treatments) the disclosure choices. In Section 5.2 we evaluate how market prices and firms’ profits differ across treatments. In Section 5.3 we analyze purchasing behavior in more detail

	FULL INFO	CHOICE-100	CHOICE-85	CHOICE-60	NO INFO
% of buyers who bought	99.6 (0.8)	99.7 (0.6)	100 (0)	100 (0)	99.7 (1.0)
Donations (offered)	61.7 (23.3)	44.5 (26.7)	31.9 (13.5)	42.9 (19.0)	15.0 (13.1)
Donations (sold)	62.1 (27.2)	43.6 (30.0)	27.3 (16.4)	40.1 (19.9)	11.5 (12.5)
% revealed (intention)	100	74.7 (12.6)	63.9 (18.7)	65.6 (15.0)	0
% revealed (realized)	100	74.7 (12.6)	55.6 (16.4)	39.2 (7.5)	0
% revealed (intention, sold only)	100	75.6 (16.6)	59.7 (23.4)	61.9 (18.3)	0
% revealed (realized, sold only)	100	75.6 (16.6)	56.3 (22.5)	48.1 (13.2)	0
Prices (offered)	43.9 (5.7)	41.7 (12.8)	39.5 (8.1)	47.1 (9.8)	43.1 (8.1)
Prices (sold)	40.6 (6.4)	36.3 (13.7)	34.4 (9.7)	44.0 (10.7)	39.5 (8.9)
Payoff buyer	79.1 (6.3)	83.5 (13.6)	85.6 (9.7)	76.0 (10.7)	80.4 (8.9)
Payoff seller	30.8 (2.8)	31.3 (5.1)	32.7 (4.4)	35.7 (3.6)	37.7 (3.9)

Notes: The table reports market averages and standard deviations (in brackets) for the different treatments.

Table 1: Summary Statistics of Main Variables of Interest

and shed light on the mechanisms behind our results.

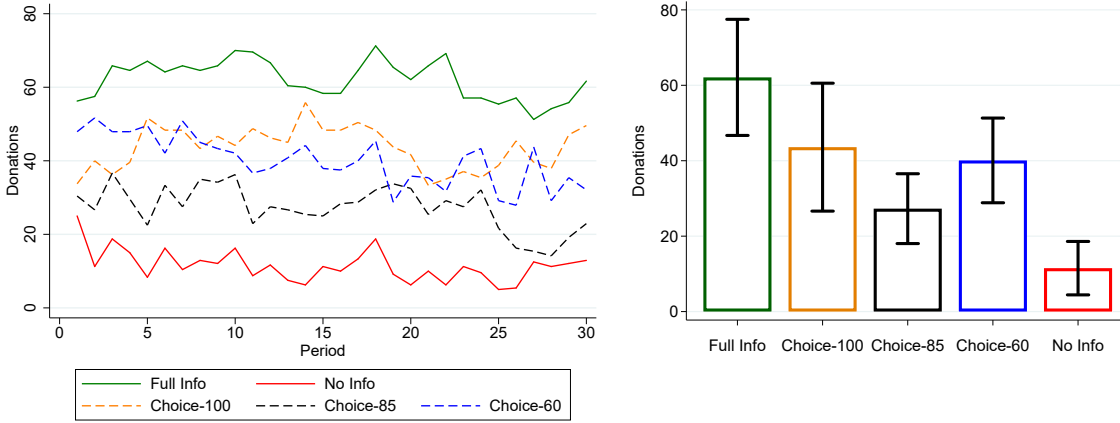
Table 1 presents summary statistics of the main variables of interest. The empirical analysis below mainly uses non-parametric tests, e.g., Mann-Whitney-U tests to test for differences across treatments, with market-level averages over the 30 periods as one independent observation.¹² Since these tests neither allow for taking into account the panel structure of the data, nor controlling for individual-level characteristics (as elicited in our post-experimental questionnaire), we complement the non-parametric analysis with parametric random-effects regressions. The regression tables can be found in the Appendix; we mention in the text those cases where the results from the latter are noteworthy.

5.1 Donations

FULL INFO vs. NO INFO

Whether or not donations are visible has a drastic influence on market outcomes. In line with Prediction 1, donations generated by a purchase in FULL INFO are almost five times larger than those in NO INFO: When donations are always observed by buyers, their purchase decisions generate donations of 62.1 points, compared to 11.5 points when donation levels are private information of an individual seller ($p < 0.001$, MWU-test). Table A1 in the Appendix shows that the corresponding treatment effect, when we control for individual

¹²The complete set of all pairwise tests for the main outcome variables can be found in Appendix A.3.



(a) Donations over Time

(b) Donations (Means and 95% CIs)

Notes: Figure (a) plots average donations per period over the thirty periods of the experiment for each of the five treatments. Figure (b) shows means and 95% confidence intervals of the donations generated on average per period across treatments.

Figure 1: Social Responsibility in the Market

buyer characteristics, is estimated as -55.7 ($p < 0.001$).

The difference between the two treatments is especially pronounced at both extremes of the possible donation levels (see Appendix Figure A1). In NO INFO, 62.0% of products bought do not generate any donation, but this is only true for 16.9% in FULL INFO. The highest donation of 100, however, is realized in only 4.2% of purchases in NO INFO, compared to 44.2% in FULL INFO. Remarkably, there is no time trend in donations as the experiment progresses (see Figure 1a).¹³

Result 1. *When the social dimension of products is visible to buyers (FULL INFO), donations are almost five times larger than when buyers do not observe it (NO INFO). This confirms Prediction 1.*

CHOICE-100 vs. FULL INFO

Average donations in CHOICE-100 are much higher than in treatment NO INFO, but appear to be smaller than in FULL INFO (see Figure 1b). In light of our predictions, especially the latter comparison is of interest. On average, in CHOICE-100, generated donations amount to 43.6 points, 29.8% less than in FULL INFO. Based on a MWU-test, this difference is however not statistically significant ($p = 0.166$). When considering the parametric

¹³See Table A1 in the Appendix for the regression that formally establishes this result by allowing for a time trend in donations.

alternative, we do estimate a significant treatment difference of -18.7 ($p = 0.044$, see Table A1).

Mirroring the comparison between NO INFO and FULL INFO, the difference, if any, between the two treatments is largely driven by behavior at the extremes of the possible donation schedule. Compared to FULL INFO, the share of accepted offers with zero donations is 7.8 percentage points higher (16.9% vs. 24.7%), whereas the share of accepted offers with the maximum donation level of 100 is 17.6 percentage points lower (44.2% vs. 26.6%).

Hence, while donations seem to be somewhat affected by making disclosure voluntary, we fail to establish unambiguous statistical significance. In order to investigate whether this is due to a potential lack of power to detect a small but significant effect of voluntary disclosure on donations, we conduct a replication of FULL INFO and CHOICE-100 with 27 markets per treatment condition (for full details of the replication, see Appendix A.5). In this replication average donations in the two treatments turn out to be very similar and we find no significant difference between the two conditions ($p = 0.736$, MWU-test). Taking these findings together, we can state the following result:

Result 2. *When sellers can choose whether to reveal the social impact of their products to buyers and disclosure is perfect (CHOICE-100), donations are not significantly different compared to the case where sellers are forced to do so (FULL INFO). This supports Prediction 2.*

Imperfect Disclosure (CHOICE-85 and CHOICE-60)

Prediction 3 tells us that when sellers decide to inform consumers about their level of donations but disclosure is imperfect, this should lower donations, compared to the FULL INFO situation. Our results indicate that this is indeed the case. Donations amount to an average of 27.3 per sold product in CHOICE-85, while in CHOICE-60, the average is 40.1. In both cases, this constitutes a significant decrease relative to donations in FULL INFO ($p = 0.003$ and $p = 0.0496$, MWU-test, respectively). Consistent with this result, the corresponding treatment effects are given by -38.0 ($p < 0.001$, CHOICE-85) and -27.0 ($p = 0.001$, CHOICE-60).

We can furthermore see a similar bimodality in donations as in the other treatments. In CHOICE-85, 11.9% of realized donations yield the highest donation, while 39.9% of accepted offers do not generate any donation at all. These shares are about halfway between NO INFO and CHOICE-100. For CHOICE-60, the share of realized donations of zero is at 36.0% and therefore comparable to behavior in CHOICE-85. Perhaps surprisingly, the share of maximal donations is substantially higher, at 23.3%. We discuss possible reasons for this in Section

5.3 below.

Result 3. *When sellers can choose to reveal the social impact of their product, but the informational content of the choice to disclose is limited, donations decrease by 56.1% in treatment CHOICE-85 and by 35.5% in treatment CHOICE-60, compared to when sellers are forced to do so and there is full transparency (FULL INFO). This supports Prediction 3.*

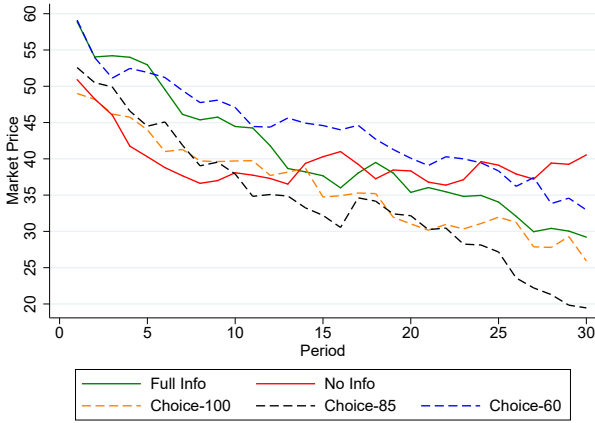
Comparing donations in CHOICE-100 and in the two imperfect disclosure treatments we do not establish a significant difference when using the non-parametric test. Yet, the regression approach in Table A1 indicates a significant treatment effect of -19.3 ($p = 0.017$) for the comparison of CHOICE-100 to CHOICE-85, but no significant difference for the comparison of CHOICE-100 to CHOICE-60 ($p = 0.985$).¹⁴ Taken together, we do not find unambiguous evidence for an effect between CHOICE-100 and the other choice treatments. Hence, uncertainty remains about how sensitive donations are to a change in disclosure probability. We will discuss this issue further in Section 5.3.

Disclosure Choice


Next, we analyze the disclosure behavior of sellers. Across the three CHOICE treatments, the share of offers where a seller decided to disclose the donation is 69.0%. The share of offers where a seller decided to disclose the donation is slightly higher in CHOICE-100 (74.7%) compared to the other two (63.9% in CHOICE-85 and 65.6% in CHOICE-60, respectively), though this difference is not statistically significant. Of course, in these two treatments, the intentions of the sellers do not correspond to the share of disclosed donations in the market. In CHOICE-85, 55.6% of offers feature a visible donation, for CHOICE-60 this number drops to 39.2%.

As discussed in Section 4, we would expect sellers to only hide their donation if they offer a donation of zero. This is indeed the case in our data. Only 9.1% of offers hide a positive donation level. Conversely, of those offers with a donation level of zero, for only 5.4% of them did the seller intend to reveal the donation to buyers. Hence, sellers, by and large, use the possibility to hide the social impact of their product when this product is socially irresponsible.

¹⁴Also, if we - based on Result 2 - pool FULL INFO with CHOICE-100 and CHOICE-85 with CHOICE-60, respectively, we can test for a difference in donations between these two pooled samples which provides a high powered test of Prediction 3. When doing so, we find a significant difference (52.8 vs. 33.7, $p = 0.036$, MWU-test).



(a) Market Prices over Time



(b) Seller Profits over Time

Notes: Figure (a) plots average market prices, i.e., accepted offers, per period over the thirty periods of the experiment for each of the five treatments. Figure (b) plots average seller earnings per period over the thirty periods of the experiment for each of the five treatments.

Figure 2: Prices and Profits in the Market

5.2 Market Prices and Profits

FULL INFO vs. NO INFO

How do these considerable differences in donations translate into market prices? Notice that buyers almost always buy a product (99.6% of cases in FULL INFO and 99.7% of cases in NO INFO). Remarkably, however, there is no significant difference in market prices between the two treatments. While in treatment FULL INFO a product is sold at an average price of 40.6, in NO INFO, buyers pay on average 39.5 ($p = 0.564$, MWU-test). In treatment NO INFO, however, market prices decline only in the first few periods while afterwards they remain almost constant over the 30 periods. In contrast, in treatment FULL INFO we observe a significant decrease in market prices by -0.902 points per period ($p < 0.001$), see Figure 2a.¹⁵

Since sellers in FULL INFO sell products which have a higher production cost due to the higher donations, they obtain lower profits than sellers in NO INFO: Sellers in FULL INFO earn on average 30.8 points compared to 37.7 points in NO INFO ($p < 0.001$, MWU-test). Controlling for observables, the corresponding treatment effect amounts to 7.2 points ($p < 0.001$), see Table A1 in the Appendix. Thus, the higher production cost due to the higher donations in FULL INFO, is almost fully borne by the sellers. In line with the result

¹⁵In NO INFO, we estimate the average per-period decrease to be -0.184 ($p = 0.222$). In both case we use RE-regression of prices as the dependent variable on a linear time trend. See Table A1 in the Appendix.

on price dynamics reported above, seller profits decrease over time in treatment FULL INFO while in NO INFO they do not.¹⁶

We note that this is not what we should observe based on Prediction 1. As discussed above, we would have expected prices to be lower in NO INFO since production cost are lower as well. The fact that prices are similar and follow a different time trend indicates that there are differences in the intensity of competition between the two treatments. We establish in Section 5.3 that purchase decisions by buyers in NO INFO hamper effective competition between sellers.

Result 4. *Market prices are not significantly different between NO INFO and FULL INFO. Sellers thus earn higher profits in treatment NO INFO. Both findings are at odds with Prediction 1.*

CHOICE-100 vs. FULL INFO

Average market prices in CHOICE-100 are 4.3 points lower than in FULL INFO. This difference is not statistically significant when using matching group averages (36.3 vs. 40.6, $p = 0.488$, MWU-test). The same result of a non-significant decrease is established when using the parametric approach (treatment effect of -4.62, $p = 0.184$). In both treatments, prices decline over the course of the experiment (see also Table A1 in the Appendix).¹⁷

As discussed in the previous subsection, hidden donations overwhelmingly mask donations of zero. An important question is whether this is reflected in the market prices. In particular, we might gain some insights as to whether buyers understand the strategic incentives of the sellers and hold unbiased beliefs. Given the evidence on limited strategic thinking in disclosure games (e.g., Jin et al., 2021; Benndorf et al., 2015; Deversi et al., 2021), it would be possible that buyers make wrong inferences about what undisclosed donations in CHOICE-100 imply. Then, they may mistakenly pay higher prices than they would for offers where they know that donations are zero.

However, our results show that in our setting buyers hold relatively accurate beliefs, correctly predicting that hidden donations mask offers with zero, or at best very small donations. This can be inferred from the fact that in FULL INFO, the average price paid for a good with a donation of zero is 26.4, compared to an average of 28.3 which is paid for a good with its donation hidden ($p = 0.681$, t-test with clustered s.e.).

¹⁶In particular, Table A1 in the Appendix shows that seller profits in Full Info decrease by 0.392 points per period ($p < 0.001$) while in No Info we estimate no significant time trend (-0.058 per period, $p = 0.409$).

¹⁷These results are confirmed in the replication study we conducted for treatments FULL INFO and CHOICE-100, see Appendix A.5 for details.

Turning to seller profits, we find that there is no meaningful difference between FULL INFO and CHOICE-100. Sellers in FULL INFO earn 30.8 points on average, compared to 31.3 points in CHOICE-100. This difference is not significant ($p = 0.908$, MWU-test). Also, in both treatments profits decline over time at a similar rate (-0.392 points per period in FULL INFO and -0.331 per period in CHOICE-100, $p = 0.397$, see Table A1). Hence, while the decrease in donations reduces the production cost for the sellers, the fact that the prices they charge are lower, especially when the donations are not revealed means that endogenizing the disclosure choice does not affect seller profits.

Result 5. *Market prices are similar between CHOICE-100 and FULL INFO. Seller profits do not differ between the two treatments. This supports Prediction 2.*

Imperfect Disclosure (CHOICE-85 and CHOICE-60)

Prediction 3 states that when the informational content of disclosure is limited, prices are predicted to be higher compared to when there is full information or perfect disclosure. Our results are at odds with this prediction. The average price paid in CHOICE-85 is 34.4 which is significantly lower than the price paid in FULL INFO, 40.6 ($p = 0.043$, MWU-test). The price paid in CHOICE-60, however, is significantly larger than the price in CHOICE-85 (44.0, $p = 0.043$, MWU-test), consistent with the intuition that lower donations lead to lower prices. Hence, there seems to be an important difference between the two treatments with respect to whether competition in the market drives down prices or not. We will return to this result further below, in particular when discussing it in relation to the differences in prices between NO INFO and FULL INFO as we argue that similar forces may be at play in both cases.

Turning to seller earnings, we find that while prices are significantly different between the two treatments, for profits this is not the case. While profits in CHOICE-60 are 3.0 points higher than in CHOICE-85, this difference fails to reach statistical significance ($p = 0.119$, MWU-test). This is explained by the slightly higher donations in CHOICE-60. Looking at the dynamics, we observe that seller earnings in both treatments decline over time, but the decrease per period is significantly stronger in CHOICE-85 than in CHOICE-60 (-0.431 vs. -0.264, $p = 0.043$, see Table A1). Moreover, profits in CHOICE-60 are higher than profits in both CHOICE-100 (35.7 vs. 31.3, $p = 0.043$, MWU-test) and FULL INFO (35.7 vs. 30.8, $p = 0.004$, MWU-test) while there is neither a difference between profits in CHOICE-85 and profits in CHOICE-100 (32.7 vs. 31.3, $p = 0.564$, MWU-test) and FULL INFO (32.7 vs. 30.8, $p = 0.225$, MWU-test).

Result 6. *Market prices in CHOICE-85 are lower than prices in FULL INFO. Prices in*

CHOICE-60 are, however, almost 10 points higher than in CHOICE-85 and not significantly different from prices in FULL INFO. This is at odds with Prediction 3. Seller profits increase when the informational content of disclosure is sufficiently limited (CHOICE-60).

Comparing profits across all five treatments, it is noteworthy that they can be ranked according to their informational content. In particular, seller profits are highest in NO INFO, followed by CHOICE-60 and CHOICE-85, which, in turn, slightly, but not significantly, outperforms the other two cases where disclosure leads to perfect information about the level of social responsibility associated with the product.¹⁸

5.3 Discussion of the Mechanisms

The results from the previous two subsections provide evidence that the disclosure (and its informativeness) of the social responsibility of products plays a crucial role both for the amount of social responsibility obtained in the market as well as for the profitability of firms competing. The goal of this subsection is to analyze the mechanisms behind the results and to highlight the implications for policy interventions, such as regulation. The discussion is guided by four questions opened up by the results.

Question 1. *Why are seller profits in NO INFO higher than in FULL INFO?*

Our two treatments where disclosure of social responsibility is exogenously determined reveal a meaningful difference both in the degree of social responsibility in the market as well as in seller profits. While the first is in line with the hypothesis, and indeed previous evidence (e.g., Pigors and Rockenbach, 2016; Etilé and Teyssier, 2016), the fact that sellers actually benefit from a regime where social responsibility is not revealed to consumers warrants closer scrutiny. In particular, while the externality is hidden from consumers, it is not absent from the market and buyers are aware that sellers had to make a decision how socially responsible to behave. Thus, buyers could try to infer donations from the only information they have, i.e., the prices posted by the sellers, but this exercise may well prove futile if sellers anticipate such behavior.

Strikingly, we find that in 30.6% of decisions, a buyer does not buy from the cheapest seller. Instead, she seems to deliberately favor a seller with a higher price, suggesting that she expects this seller to have chosen a higher donation. While it is indeed the case that in NO INFO higher prices do signal higher donations (an increase in donations by one (10

¹⁸The difference in profits between NO INFO and the other treatments is significant at the 5%-level for all treatments but CHOICE-60. Table A1 in the Appendix provides the treatment effects from the random-effects regression and confirms the findings presented in this section.

point) increment implies a 2.6-point increase in the price, $p < 0.001$, see Table A3 in the Appendix), buyers are far from perfect in inferring donations from prices (26.7% of choices are dominated by an offer with a lower price and a higher, but unobserved, donation). The fact that in FULL INFO only in 7.3% of cases buyers choose offers which are dominated by another offer with both a lower price and a higher observed donation, indicates that choosing offers with higher prices in NO INFO is not caused by erroneous behavior but stems from conscious decisions of buyers.¹⁹

It remains speculative whether buyers in NO INFO sometimes buy higher-priced products because they firmly believe that they are able to correctly predict undisclosed donations, or whether other types of reasoning, such as motivated beliefs (i.e., convincing oneself that paying a higher price is the more socially-responsible action) play an important role as well. Yet, it is clear that this buyer behavior has consequences for the competitiveness of the market: that in NO INFO in almost one third of the cases a seller can sell his product although it is not the cheapest (while dominated choices are rare in FULL INFO) yields a reduced intensity of competition in NO INFO. This explains why market prices in NO INFO display no decreasing trend, while in FULL INFO the steady decline of market prices suggests that sellers actively try to undercut each other.

At the end of the experiment, we elicited how positively subjects evaluate the charity which receives the donations. Based on these answers, we use a binary measure to classify subjects as either having a positive or a neutral/negative view.²⁰ In FULL INFO, buyers who hold a positive view about the charity generate almost three times more average donations than those who do not (74.0 vs. 26.3, $p = 0.001$, MWU-test on individual buyer averages). At the same time, in treatment NO INFO, those who think positively about the charity buy a higher priced product more than five times more often than those who do not (40.0% vs 7.6%, $p = 0.016$, MWU-test on individual buyer averages).²¹ Hence, the willingness to pay a higher price under the belief that this may generate positive donations is positively correlated with a favorable perception of the social externality.

¹⁹The full analysis of dominated choices in the different treatments can be found in Table A4 in the Appendix.

²⁰The answers were given on a scale from 1 (which corresponds to considering the charity as “very worthy of support”) to 7 (corresponding to “not at all”). We classify subjects as having a positive view of the charity if they answer either “1” or “2”. Across all participants in our sample, about 35% of participants answer “1” and another 35% answer “2”, which is why we chose this classification. There is no significant difference in the view of the charity across treatments; they are fully balanced also when comparing buyers and sellers separately.

²¹These results are robust to using “3” or “4” as a cutoff and also when using the measure as a dependent variable in a regression framework: In FULL INFO a higher value of the “charity variable” (i.e., a more negative view of the charity) has a significantly negative effect on average donations; in NO INFO a higher value significantly reduces the probability to choose a higher priced product.

Interpreting the high prevalence of choosing offers with a higher price in light of the question of (un)informative labeling and its effect on market outcomes, we predict that if there are no possibilities for sellers to inform consumers about the social responsibility of their products, e.g., because this dimension is too complex to convey, this has additional anti-competitive effects. Since consumers would have to make inferences about product characteristics that they care about but have no way of finding out, they mistakenly infer too much from prices, hampering competition and leading to higher prices.

Question 2. *How does limited transparency affect donations when disclosure is voluntary?*

Our findings that donations in CHOICE-100 come close to FULL INFO (Result 2) and that donations in CHOICE-85 and CHOICE-60 decrease compared to FULL INFO (Result 3) suggest that the decrease in donations is not caused by the fact that disclosure is voluntary, but by the limited informational content of disclosure. However, as described above, our data remains inconclusive when comparing donations in CHOICE-85 and CHOICE-60 to donations in CHOICE-100. In particular, we cannot say by how much the observed decline in social responsibility between FULL INFO and CHOICE-60 and CHOICE-85, respectively, can be attributed to the effect of making disclosure voluntary and by how much to the additional effect of a decrease in the reliability of information. One reason for this is that our initial data, as becomes clear when looking at Figure 1b again, features rather large confidence intervals, especially in CHOICE-100. Therefore, in order to shed more light on whether donations in CHOICE-100 are indeed different from FULL INFO - which would indicate that it is voluntary disclosure per se which decreases donations - we conducted a high powered replication of treatments CHOICE-100 and FULL INFO. As already described in Section 5.1 above, this replication shows that the effect of just making disclosure voluntary is small.

This observation is corroborated when estimating a random-effects regression of realized donations on a set of treatment dummy variables including all observations from the initial experiment and the replication. Since the replication was conducted in a somewhat different setting than the main experiment, we include a dummy variable indicating data coming from the replication study.²² As shown in Table A11, Column (1), we estimate a small and insignificant effect (-3.6 points, $p = 0.530$) of making disclosure voluntary on donation levels between FULL INFO and CHOICE-100. The effect of reducing the disclosure probability from 100% (CHOICE-100) to 85% (CHOICE-85) is estimated to cause a reduction in donations of 26.5 points ($p < 0.001$). The effect of moving from 100% (CHOICE-100) to 60% (CHOICE-60) is estimated to decrease donations by 16.0 points ($p = 0.039$).

²²As described in more detail in Appendix A.5, the main (and unavoidable) difference was that due to the COVID-19 pandemic, the replication study had to be conducted online.

Hence, while we acknowledge that some uncertainties remain regarding the interpretation of the findings of the three CHOICE treatments, we view these results as suggestive evidence that the observed decline in social responsibility between FULL INFO and CHOICE-60 and CHOICE-85, respectively, is to a large extent due to the effect of a decrease in the reliability of information, while simply making disclosure voluntary does not affect social responsibility much.

Question 3. *What differences exist between CHOICE-85 and CHOICE-60 and why?*

A more detailed analysis of the two imperfect disclosure treatments can provide further insights into how conditions which make it hard for firms to provide easily-interpretable and unambiguous information about the social responsibility of their products to consumers, affect market outcomes. As discussed above, prices in CHOICE-60 are significantly higher compared to CHOICE-85 (Result 6) and sellers earn significantly more in CHOICE-60 than in treatments CHOICE-100 and FULL INFO. Seller earnings in CHOICE-85, however, are not significantly different from those in FULL INFO and CHOICE-100, respectively

To make sense of these results, note that given the disclosure decisions of the sellers and the imperfect implementation probability of 60%, only for 39% of offers in CHOICE-60 do consumers know the donation level. In these cases, the decision which product to buy bears some resemblance to the NO INFO treatment in the sense that buyers might engage in inferences about the donation level based on the prices. Corroborating our earlier finding for the NO INFO treatment, we now find that in 13.8% of decisions in CHOICE-60, a buyer buys a product with a hidden donation which is not the cheapest among those with a hidden donation. In comparison, in CHOICE-85 this only happens in 4.3% of cases.

This suggests that behavior in CHOICE-60 is similar to NO INFO in the sense that buyers try to predict donation levels for those donations they do not observe. The difference between the two treatments is that in CHOICE-60 buyers do sometimes get information about seller behavior. This might allow sellers to create a reputation for being a socially responsible seller (remember that sellers' identities are known to buyers and do not change over time). This may explain why there are relatively many donations of 100 in this treatment, compared to CHOICE-85 (23.3% vs. 11.9%). At the same time, having developed a reputation for being socially responsible could create leeway for sellers to charge higher prices, consistent with what we observe in the data. We discuss the role that reputation plays in the CHOICE treatments in more detail below.

Question 4. *What role does reputation play in the CHOICE-treatments?*

In the three CHOICE-treatments a buyer will regularly be confronted with offers with

	(1)	(2)	(3)	(4)
Weight on Buyer Earnings (β)	0.133*** (0.020)	0.143*** (0.021)	0.136*** (0.020)	0.146*** (0.021)
Weight on Donations (θ^b)	0.019*** (0.003)	0.022*** (0.003)	0.028*** (0.009)	0.029*** (0.008)
Donations (θ^b) \times CHOICE-85			-0.012 (0.010)	-0.012 (0.010)
Donations (θ^b) \times CHOICE-60			-0.010 (0.009)	-0.008 (0.009)
Hidden Donation Dummy	-0.133 (0.157)	-0.081 (0.163)	0.060 (0.280)	-0.027 (0.277)
Hidden Donation Dummy \times CHOICE-85			-0.283 (0.347)	-0.063 (0.338)
Hidden Donation Dummy \times CHOICE-60			-0.350 (0.387)	-0.196 (0.406)
Seller Share	0.629*** (0.221)	0.470** (0.213)	0.159 (0.216)	0.059 (0.247)
Seller Share \times CHOICE-85			0.311 (0.404)	0.349 (0.518)
Seller Share \times CHOICE-60			1.036** (0.456)	0.814* (0.441)
Case-specific individual variables	No	Yes	No	Yes
Observations	7302	7302	7302	7302
Cases	2086	2086	2086	2086
Cluster	36	36	36	36

Notes: The table shows estimates from Mc Fadden's conditional logit choice models. We assume that a buyer's utility from buying a product for a price p and a donation d is given by $u = \beta(120 - p) + \theta^b d$. $\beta > 0$ denotes the weight a consumer puts on her own earnings, while θ^b captures the extent to which a consumer cares about the donations to unicef. The seller share indicates the share of offers a buyer accepted from the respective seller in the past. All models contain data from CHOICE-100, CHOICE-85 and CHOICE-60, with CHOICE-100 serving as the reference category. The treatment variables are binary. Columns (2) and (4) also contain case specific individual variables including period-specific dummy variables and characteristics of the buyer. Standard errors (clustered at the market level) in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2: Conditional Logit Choice Model for CHOICE-Treatments

hidden donations. But, since sellers are identified through letters A to D in the experiment, buyers have the possibility to track seller behavior across periods. This might allow them to make inferences about hidden donations based on sellers' past behavior. At the same time, sellers might therefore be able to develop a reputation for being, for example, a socially responsible seller. Buyers might then be more likely to also buy from such a seller even if the donation is hidden in some of the periods.

In order to analyze whether reputation plays a role in our markets, we estimate, similar to Bartling et al. (2015), a conditional logit choice model. We assume that a buyer's utility from buying a product for a price p and a donation d is given by $u = \beta(120 - p) + \theta^b d$. Here, $\beta > 0$ denotes the weight a consumer puts on her own monetary payoff, while θ^b captures the extent to which a consumer cares about the level of social responsibility of a product, i.e., the donations generated by purchasing it. When estimating the coefficients β and θ^b we treat donations which are observable to the buyer differently from those which are not and capture the latter via a dummy variable equal to one whenever a donation is not visible to the buyer.

Our main interest, however, lies in using this conditional choice model for quantifying a potential effect of reputation. To do that, we include an additional variable into this model which captures how often a buyer previously bought from a specific seller. More precisely, for each alternative available to a buyer we calculate the share of offers she accepted from this specific seller in the past. Hence, controlling for the price and donation of a given offer, we check whether a buyer is more likely to buy from a seller he already bought more from in the past.

Table 2 presents the results from estimating such a model. We first note that we estimate, as should be expected, a positive value for β as well as θ^b , with the ratio of the two being between five and seven, depending on the specification, indicating that buyers are willing to give up one point of their own monetary payoff to increase donations by five to seven points. We then note the share of offers previously bought from this seller has a positive effect on the utility associated with a specific alternative (see columns (1) and (2) in Table 2). When allowing for this effect to vary across treatments, we find a significant effect only in the CHOICE-60 treatment. To quantify this effect, a ten percentage points increase in the share of offers bought from a seller is comparable in magnitude to a price increase of about one point. In CHOICE-85 and CHOICE-100, the effect is not significant.

Hence, corroborating our results on the dominated choices presented above, we do indeed find evidence that some form of reputation plays a role in those CHOICE-treatments where disclosure is imperfect. However, we note that our experiment was not designed to specifically

analyze the role of reputation. It would be an interesting avenue for further research to analyze the role of reputation more thoroughly.

6 Implications and Concluding Remarks

In this paper, we have analyzed the role of social responsibility (SR) in markets. We show that disclosure of SR information has large and positive effects compared to a regime where producers are unable to disclose the social responsibility of their products to consumers (NO INFO). We find that when voluntary disclosure provides unambiguous information for customers (CHOICE-100), the level of social responsibility comes close to the outcome under full and externally enforced disclosure (FULL INFO). Yet, in the cases where producers cannot be sure that the provided SR information has full informational content for consumers (CHOICE-85 and CHOICE-60), SR in production is significantly lower, compared to FULL INFO. Yet, since our data does not provide an unambiguous statistical difference between the three choice treatments, the sensitivity of social responsibility to changes in the disclosure probability cannot be precisely determined.

We believe that for policy makers and regulators interested in influencing the level of SR in the market, several implications of our results are noteworthy. First, we have established that the more opaque the disclosure of SR attributes, the larger the firms' profits. Hence, one should expect firms to oppose, on the one hand, those kinds of regulations which enforce disclosure of such attributes, but also, on the other hand, initiatives which facilitate comparisons of products along this dimension.

On a more positive note, our results show that the benefits from full transparency and comparability can be large compared to a situation with no information, even if one relies on voluntary disclosure by firms. Especially as too heavy-handed regulation may be difficult to implement and enforce, voluntary disclosure can go a long way. Our results indicate that competition on the SR dimension is sufficiently strong to deliver both high disclosure rates as well as high levels of SR.

Importantly, however, our results indicate that the above effects of voluntary disclosure are limited in markets where the producer faces uncertainty about whether disclosed SR information reaches the consumers. We studied a situation in which the SR of the product - when communicated to consumers - was unambiguously clear. The overwhelming majority of subjects knew *unicef* and rated it highly positive. Yet it seems reasonable to assume that the effects of voluntary disclosure get weaker when SR information is less meaningful, either because the provided information captures only a part of a complex and opaque production

process or a multiplicity of labels distorts the informational content and limits comparability. These results thus provide empirical support for theoretical models which posit that consumer confusion about the informational content of labels decrease the incentive for firms to disclose these attributes (e.g., Harbaugh et al., 2011). They also suggest that an important role for policy makers lies in making it easier for firms to communicate to consumers how to assess the SR of products, for example through reliable and easy-to-interpret information.

References

- BARTLING, B., E. FEHR, AND Y. ÖZDEMİR (2020a): “Does market interaction erode moral values?” *The Review of Economics and Statistics*, 1–32.
- BARTLING, B., V. VALERO, AND R. WEBER (2019): “On the scope of externalities in experimental markets,” *Experimental Economics*, 22, 610–624.
- BARTLING, B., V. VALERO, AND R. A. WEBER (2020b): “The Causal Effect of Income Growth on Consumer Social Responsibility,” Tech. rep., University of Zurich.
- BARTLING, B., R. A. WEBER, AND L. YAO (2015): “Do markets erode social responsibility?” *The Quarterly Journal of Economics*, 130, 219–266.
- BENNDORF, V., D. KÜBLER, AND H.-T. NORMANN (2015): “Privacy concerns, voluntary disclosure of information, and unraveling: An experiment,” *European Economic Review*, 75, 43–59.
- BUND (2019): “Mit Brief und (Bio-)Siegel: Welche Kennzeichnung von Lebensmitteln ist empfehlenswert?” <http://www.bund.net/massentierhaltung/haltungskennzeichnung/bio-siegel/>, accessed: 2019-05-16.
- DABNEY, C. (2017): “The Truth About Coffee Packaging Claims,” <http://www.consumerreports.org/coffee/the-truth-about-coffee-packaging-claims/>, accessed: 2019-05-16.
- DEVERSI, M., A. ISPANO, AND P. SCHWARDMANN (2021): “Spin Doctors: An Experiment on Vague Disclosure,” Working Paper.
- DUCH, M. L., M. R. GROSSMANN, AND T. LAUER (2020): “z-Tree unleashed: A novel client-integrating architecture for conducting z-Tree experiments over the Internet,” *Journal of Behavioral and Experimental Finance*, 28, 100400.
- ELFENBEIN, D. W. AND B. MCMANUS (2010): “A greater price for a greater good? Evidence that consumers pay more for charity-linked products,” *American Economic Journal: Economic Policy*, 2, 28–60.
- ENGELMANN, D., J. FRIEDRICHSEN, AND D. KÜBLER (2018): “Fairness in markets and market experiments,” WZB Discussion Paper.
- ETILÉ, F. AND S. TEYSSIER (2016): “Signaling Corporate Social Responsibility: Third-Party Certification versus Brands,” *Scandinavian Journal of Economics*, 118, 397–432.
- FALK, A. AND N. SZECH (2013): “Morals and markets,” *Science*, 340, 707–711.
- FAUL, F., E. ERDFELDER, A. BUCHNER, AND A.-G. LANG (2009): “Statistical power analyses using G* Power 3.1: Tests for correlation and regression analyses,” *Behavior research methods*, 41, 1149–1160.
- FEICHT, R., V. GRIMM, AND M. SEEBAUER (2016): “An experimental study of corporate social responsibility through charitable giving in Bertrand markets,” *Journal of Economic Behavior & Organization*, 124, 88–101.
- FIBL (2017): Data on organic agriculture 2001-2017. The Statistics.FiBL.org website maintained by the Research Institute of Organic Agriculture (FiBL), Frick, Switzerland. Data are available at <http://statistics.fibl.org/world.html>, accessed: 2019-05-16.
- FISCHBACHER, U. (2007): “z-Tree: Zurich toolbox for ready-made economic experiments,” *Experimental Economics*, 10, 171–178.

- GREINER, B. (2015): “Subject pool recruitment procedures: organizing experiments with ORSEE,” *Journal of the Economic Science Association*, 1, 114–125.
- HAINMUELLER, J., M. J. HISCOX, AND S. SEQUEIRA (2015): “Consumer demand for fair trade: Evidence from a multistore field experiment,” *Review of Economics and Statistics*, 97, 242–256.
- HARBAUGH, R., J. W. MAXWELL, AND B. ROUSSILLON (2011): “Label confusion: The Groucho effect of uncertain standards,” *Management Science*, 57, 1512–1527.
- HENZE, B., F. SCHUETT, AND J. P. SLUIJS (2015): “Transparency in markets for experience goods: experimental evidence,” *Economic Inquiry*, 53, 640–659.
- HOLT, C. AND R. SHERMAN (1990): “Advertising and Product Quality in Posted-Offer Experiments,” *Economic Inquiry*, 28, 39–56.
- IRLENBUSCH, B. AND M. C. VILLEVAL (2015): “Behavioral ethics: how psychology influenced economics and how economics might inform psychology?” *Current Opinion in Psychology*, 6, 87–92.
- ISPANO, A. AND P. SCHWARDMANN (2018): “Competition over Cursed Consumers,” Working Paper.
- JIN, G. Z., M. LUCA, AND D. MARTIN (2021): “Is No News (Perceived As) Bad News? An Experimental Investigation of Information Disclosure,” .
- KIRCHLER, M., J. HUBER, M. STEFAN, AND M. SUTTER (2016): “Market design and moral behavior,” *Management Science*, 62, 2615–2625.
- OTA (2018): Organic Trade Association, Press Release on Organic Industry Survey 2017, May 18, 2018, <http://ota.com/news/press-releases/20236>, accessed: 2019-05-16.
- PIGORS, M. AND B. ROCKENBACH (2016): “Consumer social responsibility,” *Management Science*, 62, 3123–3137.
- RODE, J., R. M. HOGARTH, AND M. LE MENESTREL (2008): “Ethical differentiation and market behavior: An experimental approach,” *Journal of Economic Behavior & Organization*, 66, 265–280.
- SUTTER, M., J. HUBER, M. KIRCHLER, M. STEFAN, AND M. WALZL (2020): “Where to look for the morals in markets?” *Experimental Economics*, 23, 30–52.

Online Appendix

A Additional Analyses

A.1 Regressions

	Donations		Price		Seller Earnings	
	(1)	(2)	(3)	(4)	(5)	(6)
CHOICE-100	-18.735** (9.311)	-21.661** (10.553)	-4.615 (3.475)	-7.666 (5.168)	0.507 (1.586)	-0.441 (2.320)
CHOICE-85	-37.987*** (8.070)	-36.004*** (9.583)	-7.906** (3.276)	-6.409* (3.821)	1.960 (1.457)	2.557 (2.113)
CHOICE-60	-27.023*** (8.127)	-21.855** (8.970)	1.831 (3.179)	-1.151 (3.529)	4.826*** (1.282)	2.844* (1.648)
NO INFO	-55.719*** (7.936)	-56.198*** (8.136)	-2.436 (3.253)	-13.546*** (3.761)	7.148*** (1.389)	1.970 (1.868)
Period		-0.252 (0.239)		-0.902*** (0.120)		-0.392*** (0.044)
CHOICE-100 × Period		0.189 (0.496)		0.198 (0.217)		0.061 (0.072)
CHOICE-85 × Period		-0.128 (0.425)		-0.096 (0.183)		-0.039 (0.081)
CHOICE-60 × Period		-0.333 (0.302)		0.193 (0.150)		0.128** (0.063)
NO INFO × Period		0.031 (0.247)		0.718*** (0.193)		0.334*** (0.082)
Observations	3600	3600	3593	3593	7200	7200
Cluster	60	60	60	60	60	60

Notes: The table shows the results from random-effects regressions on dummy variables indicating the treatment and the individual-specific control variables described in Section A.1.1. Columns (1) (3) and (5) also contain Period-specific dummy variables. Standard errors are clustered by market. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A1: Treatment Effects for Main Outcomes

	Offered Donations		Buyer Earnings	
	(1)	(2)	(3)	(4)
CHOICE-100	-16.395*	-20.202**	4.586	7.788
	(8.850)	(9.035)	(3.445)	(5.131)
CHOICE-85	-29.835***	-27.619***	8.079**	6.607*
	(7.822)	(7.639)	(3.258)	(3.747)
CHOICE-60	-19.021**	-12.284	-1.703	1.307
	(8.297)	(7.878)	(3.169)	(3.473)
NO INFO	-46.812***	-46.621***	2.363	13.669***
	(7.205)	(6.070)	(3.282)	(3.685)
Period		-0.021		0.903***
		(0.202)		(0.121)
CHOICE-100 × Period		0.246		-0.207
		(0.395)		(0.220)
CHOICE-85 × Period		-0.143		0.095
		(0.350)		(0.184)
CHOICE-60 × Period		-0.435*		-0.194
		(0.264)		(0.151)
NO INFO × Period		-0.012		-0.729***
		(0.213)		(0.197)
Observations	7200	7200	3600	3600
Cluster	60	60	60	60

Notes: The table shows the results from random-effects regressions on dummy variables indicating the treatment and the individual-specific control variables described in Section A.1.1. Columns (1) (3) and (5) also contain Period-specific dummy variables. Standard errors are clustered by market. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A2: Treatment Effects for Additional Outcomes

	(1)	(2)	(3)	(4)	(5)
	FULL INFO	CHOICE-100	CHOICE-85	CHOICE-60	NO INFO
Offered Donations	0.234*** (0.024)	0.254*** (0.032)	0.249*** (0.025)	0.181*** (0.031)	0.260*** (0.054)
Observations	1440	1440	1440	1440	1440
Cluster	12	12	12	12	12

Notes: The table shows the results from random-effects regressions of the offer on the donation level associated with the product for the five treatments. The regressions also include a set of individual-specific control variables. Standard errors are clustered by market. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A3: Correlations between Prices and Donations

A.1.1 Additional Control Variables

The following variables are used as additional control variables in the regressions:

- *age*: Participant’s age
- *Degree*: Dummy variable indicating the degree the participant plans to obtain from the course currently enrolled in (Bachelor= 1, Master= 2, other= 0).
- *Female*: Dummy variable indicating the participant’s gender.
- *German*: Dummy variable indicating whether the participant’s nationality is German, or not.
- *unicef*: Do you, in principle, consider unicef an organization worthy of support? (scale from 1 to 7 where 1=“yes, very” and 7=“no, not at all”. 1 is omitted category)
- *donation frequency*: Have you donated money to charity in the last twelve months? (where 1=“yes, weekly”, 2=“yes, monthly”, 3=“yes, from time to time”, 4=“yes, at least once”, and 5=“no”. 5 is omitted category)
- *donation size*: If yes, how much did you donate approximately in the last year? (in Euro)

A.2 Dominated Choices

	FULL INFO	CHOICE-100	CHOICE-85	CHOICE-60	NO INFO
Dominated Choices	52/717 (0.0725)	55/518 (0.1062)	11/355 (0.0310)	6/210 (0.0286)	—
Price-dominated Choices	—	7/117 (0.0598)	11/257 (0.0428)	48/347 (0.1383)	220/718 (0.3064)

Notes: The table shows dominated and price-dominated choices, separately by treatment. For the first row (dominated choices), we consider all cases where a buyer had at least two product with visible donations to choose from and chose one of these. A choice is classified as dominated if for the chosen product (p_i, d_i) , there is another product (p_j, d_j) with a visible donation and $p_i \geq p_j$ and $d_i \leq d_j$ with at least one inequality strict. For the second row (price-dominated choices), we consider all cases where a buyer had at least two products with hidden donations to choose from and chose one of these. A choice is classified as price-dominated if for a chosen product (p_i, \cdot) , there is another product (p_j, \cdot) with a hidden donation and $p_i > p_j$. In both cases, the first number denotes the number of (price-)dominated choices among all applicable cases. The number in brackets denotes the share of (price-)dominated choices.

Table A4: Dominated Choices by Treatment

A.3 Non-parametric Tests

	FULL INFO	CHOICE-100	CHOICE-85	CHOICE-60	NO INFO
Donations (offered)	61.72 (23.34)	44.50 (26.66)	31.88 (13.53)	42.85 (18.98)	14.97 (13.13)
Donations (sold)	62.10 (27.21)	43.60 (29.98)	27.29 (16.36)	40.08 (19.85)	11.51 (12.52)

	p-value (Donations, offered)	p-value (Donations, sold)
FULL INFO vs. CHOICE-100	0.119	0.166
FULL INFO vs. CHOICE-85	0.002	0.003
FULL INFO vs. CHOICE-60	0.053	0.0496
FULL INFO vs. NO INFO	<0.001	<0.001
CHOICE-100 vs. CHOICE-85	0.326	0.225
CHOICE-100 vs. CHOICE-60	0.885	0.773
CHOICE-100 vs. NO INFO	0.003	0.001
CHOICE-85 vs. CHOICE-60	0.149	0.126
CHOICE-85 vs. NO INFO	0.011	0.013
CHOICE-60 vs. NO INFO	0.001	<0.001

Notes: The table reports p-values from MWU-tests for treatment differences. The left column presents results for donations for all offered products. The right column uses donations associated with sold products only. The level of independent observations are matching group averages. The top part of the table presents, for completeness, the averages, as presented in Table 1 in the main text.

Table A5: MWU-tests for treatment differences: donations

	FULL INFO	CHOICE-100	CHOICE-85	CHOICE-60	NO INFO
Prices (offered)	43.85 (5.66)	41.65 (12.75)	39.51 (8.08)	47.11 (9.80)	43.07 (8.08)
Prices (sold)	40.64 (6.44)	36.28 (13.68)	34.41 (9.72)	44.03 (10.70)	39.48 (8.92)

	p-value (Prices, offered)	p-value (Prices, sold)
FULL INFO vs. CHOICE-100	0.488	1
FULL INFO vs. CHOICE-85	0.043	0.0496
FULL INFO vs. CHOICE-60	0.299	0.204
FULL INFO vs. NO INFO	0.564	0.488
CHOICE-100 vs. CHOICE-85	0.773	0.564
CHOICE-100 vs. CHOICE-60	0.225	0.419
CHOICE-100 vs. NO INFO	0.624	0.954
CHOICE-85 vs. CHOICE-60	0.043	0.024
CHOICE-85 vs. NO INFO	0.166	0.248
CHOICE-60 vs. NO INFO	0.166	0.094

Notes: The table reports p-values from MWU-tests for treatment differences. The left column presents results for all offers. The right column uses prices of sold products only. The level of independent observations are matching group averages. The top part of the table presents, for completeness, the averages, as presented in Table 1 in the main text.

Table A6: MWU-tests for treatment differences: prices


	FULL INFO	CHOICE-100	CHOICE-85	CHOICE-60	NO INFO
Payoff buyer	79.10 (6.33)	83.53 (13.60)	85.59 (9.71)	75.97 (10.70)	80.35 (8.93)
Payoff seller	30.79 (2.83)	31.25 (5.10)	32.71 (4.43)	35.72 (3.64)	37.68 (3.91)

	p-value (Payoff buyer)	p-value (Payoff seller)
FULL INFO vs. CHOICE-100	0.488	0.908
FULL INFO vs. CHOICE-85	0.021	0.225
FULL INFO vs. CHOICE-60	0.326	0.004
FULL INFO vs. NO INFO	0.488	<0.001
CHOICE-100 vs. CHOICE-85	0.686	0.564
CHOICE-100 vs. CHOICE-60	0.248	0.043
CHOICE-100 vs. NO INFO	0.624	0.009
CHOICE-85 vs. CHOICE-60	0.043	0.119
CHOICE-85 vs. NO INFO	0.149	0.013
CHOICE-60 vs. NO INFO	0.166	0.204

Notes: The table reports p-values from MWU-tests for treatment differences. The left column presents results for buyer payoffs, the right column for seller payoffs.. The level of independent observations are matching group averages. The top part of the table presents, for completeness, the averages, as presented in Table 1 in the main text.


Table A7: MWU-tests for treatment differences: payoffs

A.4 Additional Figures



Notes: Figure plots a histogram of the donations generated by the purchase decisions of the buyers for each of the five treatments. In the Choice treatments, shaded bars denote hidden donations.

Figure A1: Histograms of Donations



(a) Offers (incl. rejected) over Time

(b) Buyer earnings over Time

Notes: Figure (a) plots average offers (all offers, accepted and rejected) per period over the thirty periods of the experiment for each of the five treatments. Figure (b) plots average buyer earnings per period over the thirty periods of the experiment for each of the five treatments.

Figure A2: Offers and Buyer Earnings in the Market

A.5 Replication Study

In March 2021, we ran a replication study of FULL INFO and CHOICE-100 to address the inconclusive results regarding differences between these two conditions. As the Cologne Laboratory for Economic Research (CLER) has been closed throughout the COVID-19 pandemic, the experiment was conducted online via zTree unleashed (Duch et al., 2020), an online architecture for running zTree experiments over the internet. Like in the initial experiment, participants were recruited from the subject pool of the CLER via ORSEE (Greiner, 2015).

We sampled 28 markets per treatment condition, which provides 80% power to detect an effect of the size observed on donations in the initial experiment (0.65 sd) at the 5% significance level using Mann-Whitney-U tests. The power calculations were performed in G*Power 3.1 (Faul et al., 2009) assuming a logistic distribution as the parent distribution and performing a one sided-test. The one sided-test is justified given the directional hypothesis from the initial experiment. The achieved power is to some degree sensitive to the assumption of the parent distribution: assuming a normal distribution yields 75% power in the above specification, while a Laplace distribution yields 90% power. Two of the 56 markets had to be terminated due to technical difficulties, one in each treatment condition, so that 27 markets per treatment remain for analysis, which affects power only marginally. In the initial experiment with 12 markets per treatment condition we were powered to detect a minimal effect size of about one standard deviation (power=0.80%, $\alpha = 5\%$, MWU test, one-sided).

As performing an exact replication was not possible due to the pandemic, there are some unavoidable differences in the experimental procedures between the initial experiment and the replication, which could limit the comparability across the two experiments. Most importantly, subjects participated in the online experiment from home, so that the degree of social interactions between subjects was much lower. During the online experiment subjects joined a Zoom meeting in which the instructions for the experiment were provided and technical difficulties could be resolved. The instructions for the online experiment differed from the instructions in the laboratory only in the introductory paragraph (see Appendix B.3). Subjects did not have to turn on their camera while in the Zoom meeting. In that way the online experiment was arguably more anonymous compared to the laboratory. Payments were made via PayPal instead of cash payments used in the laboratory experiment. On average, the online experiment also took around 30 minutes more time than the laboratory experiment.

To conduct the online replication, we also had to make a number of changes to the initial zTree program. First of all, we had to add screens to collect the PayPal addresses of sub-

jects. Secondly, we had to implement a procedure to handle attrition of subjects. In case the internet connection to a subject was lost, they had 3 minutes time to reconnect to the experiment. Otherwise, the experiment was terminated for all subjects in the respective market. As mentioned above, this happened only in two markets, one per treatment condition. Last, we had to make a number of minor graphical adjustments.

A.5.1 Results - Replication Only

Summarizing our results, we find that the level of donations, prices and earnings do not differ significantly between FULL INFO and CHOICE-100. Hence, the online replication provides unambiguous evidence for Prediction 2. In what follows we describe the main findings in more detail.

Average donations in FULL INFO amount to 41.25, compared to 42.54 in CHOICE-100 ($p = 0.959$, MWU-test, see Table A9 for a complete overview of all tests). Figure A3a shows that the histograms of the realized donations and the level of donations over time exhibit no noticeable differences between the two treatments.

The prices and profits over time depicted in Figure A4a and A4b show a similar trend, just that both outcomes tend to be slightly higher in CHOICE-100 in the first half of the experiment. In the second half of the experiment the prices and profits in the two treatments converge. Average prices in CHOICE-100 are given by 35.19, while average prices in FULL INFO are 31.83 ($p = 0.239$).

Also following our analysis of the main experiment, we compare the quality of decisions of subjects between the online experiment and the laboratory experiment. When taking the share of dominated choices as a proxy for the quality of decisions, the decision of subjects in the online experiment are clearly not worse than in the laboratory (compare Table A10 with Table A4). Hence, we find no evidence that subjects, for example, paid less attention in the online experiment compared to the laboratory experiment.

	FULL INFO	CHOICE-100
% of buyers who bought	99.9 (0.4)	99.5 (1.0)
Donations (offered)	43.5 (23.0)	42.6 (25.2)
Donations (sold)	41.2 (26.0)	42.5 (29.6)
% revealed	100	71.5 (18.3)
% revealed (sold only)	100	70.5 (23.4)
Prices (offered)	36.9 (10.1)	39.9 (8.9)
Prices (sold)	31.8 (11.3)	35.2 (11.2)
Payoff buyer	88.1 (11.3)	84.5 (11.1)
Payoff seller	29.2 (3.6)	30.8 (4.9)


Notes: The table reports market averages and standard deviations (in brackets) for the different treatments.

Table A8: Summary Statistics of Main Variables of Interest (Replication)

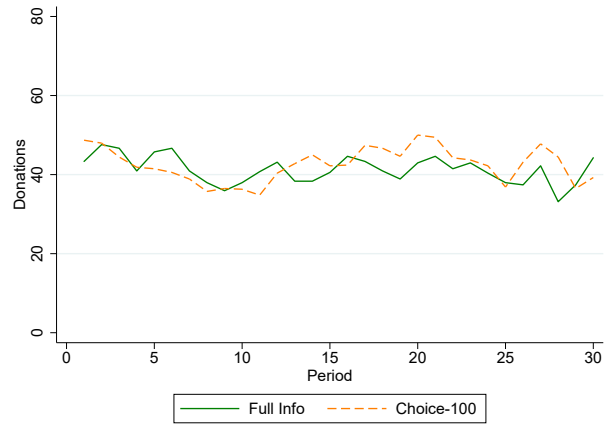
	p-value (FULL INFO vs. CHOICE-100)
Donations (offered)	0.736
Donations (sold)	0.959
Prices (offered)	0.283
Prices (sold)	0.239
Payoff (buyer)	0.229
Payoff (seller)	0.328

Notes: The table reports p-values from MWU-tests for treatment differences. The level of independent observations are matching group averages.

Table A9: MWU-tests for treatment differences (Replication)




(a) Histograms of Realized Donations



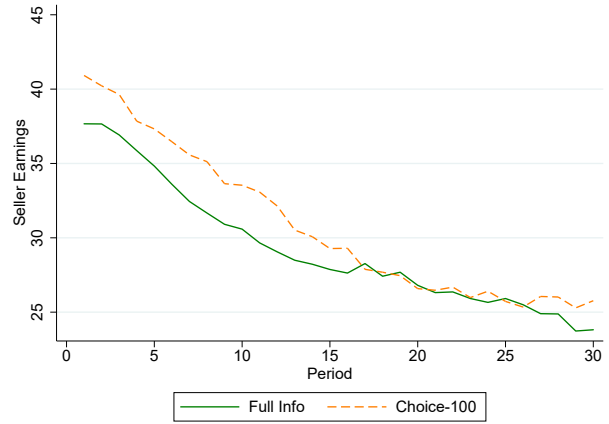
(b) Donations over Time

Notes: Figure (a) plots a histogram of the donations generated by the purchase decisions of the buyers for each of the five treatments. In the CHOICE treatments, shaded bars denote hidden donations. Figure (b) plots average donations per period over the thirty periods of the experiment for each of the five treatments.

Figure A3: Social Responsibility in the Market (Replication)



(a) Market Prices over Time



(b) Seller Profits over Time

Notes: Figure (a) plots average market prices, i.e., accepted offers, per period over the thirty periods of the experiment for each of the five treatments. Figure (b) plots average seller earnings per period over the thirty periods of the experiment for each of the five treatments.

Figure A4: Prices and Profits in the Market (Replication)

	FULL INFO	CHOICE-100
Dominated Choices	52/1,618 (0.0321)	69/1,079 (0.0639)
Price-dominated Choices	—	18/320 (0.0563)

Notes: The table shows dominated and price-dominated choices, separately by treatment. For the first row (dominated choices), we consider all cases where a buyer had at least two product with visible donations to choose from and chose one of these. A choice is classified as dominated if for the chosen product (p_i, d_i) , there is another product (p_j, d_j) with a visible donation and $p_i \geq p_j$ and $d_i \leq d_j$ with at least one inequality strict. For the second row (price-dominated choices), we consider all cases where a buyer had at least two products with hidden donations to choose from and chose one of these. A choice is classified as price-dominated if for a chosen product (p_i, \cdot) , there is another product (p_j, \cdot) with a hidden donation and $p_i > p_j$. In both cases, the first number denotes the number of (price-)dominated choices among all applicable cases. The number in brackets denotes the share of (price-)dominated choices.

Table A10: Dominated Choices by Treatment (Replication)

A.5.2 Results - Combining Replication and Main Experiment

	Donations	Price	Seller Earnings	Offered Donations	Buyer Earnings
	(1)	(2)	(3)	(4)	(5)
Choice-100	-3.557 (5.659)	0.320 (2.140)	1.278 (0.910)	-6.388 (5.272)	-0.433 (2.142)
Choice-85	-30.049*** (7.514)	-5.267 (3.272)	2.306 (1.454)	-25.281*** (6.595)	5.353 (3.259)
Choice-60	-19.541*** (7.571)	4.166 (3.312)	5.147*** (1.271)	-13.874* (7.310)	-4.125 (3.311)
No Info	-47.586*** (6.849)	0.295 (3.180)	7.538*** (1.319)	-41.930*** (6.132)	-0.457 (3.209)
replication	-19.455*** (6.875)	-5.094* (2.794)	-1.316 (1.217)	-7.283 (6.227)	4.845* (2.776)
Observations	6840	6823	13680	13680	6840
Cluster	114	114	114	114	114

Notes: The table shows the results from random-effects regressions on dummy variables indicating the treatment, a dummy indicating data collected in the replication and individual-specific control variables described in Section A.1.1. Each regression also contains Period-specific dummy variables. Standard errors are clustered by market. The regressions include all observations from the initial experiment and the replication. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A11: Treatment Effects on Donations - Main Experiment and Replication

In Table A11 we present the results of a regression using both the data from our main experiment as well as from the replication study. In the main text, we discussed in detail the results of Column (1) of Table A11, that is the treatment effects regarding donations. In the following we show that for the treatment effects on prices (Column (2)) and seller earnings (Column (3)) the results are very similar compared to the regression results in Table A1, i.e., using only the data from the main experiment.

In particular, to highlight some noteworthy cases:

- Market prices are significantly higher in CHOICE-60 than in CHOICE-85 (treatment

effect: 9.43, $p = 0.011$).

- Seller earnings in CHOICE-60 are significantly higher than in FULL INFO (treatment effect: 5.15, $p < 0.001$) and than in CHOICE-100 (treatment effect: 3.87, $p = 0.006$).
- Seller earnings in CHOICE-85 are neither significantly different from seller earnings in FULL INFO (treatment effect: 2.31, $p = 0.112$) nor than seller earnings in CHOICE-100 (treatment effect: 1.03, $p = 0.510$).

A.6 Additional Treatment

As discussed in the main text in footnote 9, we implemented one additional treatment, similar to CHOICE-100 as a robustness check. In this treatment, which we call CHOICE-100-COSTLY, we implement a donation schedule $\tilde{c}(d)$ which makes all *positive* donations 3 points cheaper, while no donation remains costless, i.e.,:

$$\tilde{c}(d) = \begin{cases} c(d) - 3 & \text{if } d > 0 \\ 0 & \text{if } d = 0 \end{cases}$$

At the same time, disclosing the donation associated with a product no longer is costless, sellers incur a cost of 3 points.

This implies that whenever a seller decides to disclose and chooses a non-zero donation, the cost of this action is the same in CHOICE-100 and CHOICE-100-COSTLY. The only difference occurs in the case where a seller decides to choose a positive donation but not to disclose. This is cheaper (by 3 points) in CHOICE-100-COSTLY.

Our motivation for running this treatment was as follows: Compared to CHOICE-100, a seller in CHOICE-100-COSTLY who cares sufficiently strongly about offering socially responsible products might be willing to offer products with high donations but wants to save on the disclosure cost. She might, for example, rather reduce the price by three points in order to make the product more attractive to a buyer. If buyers hold the belief that such sellers exist, they might be more willing to buy products with undisclosed donations. This, in turn, increases the incentives for all sellers to offer products with hidden donations, exploiting the belief of the buyers by choosing donations of zero.

We thus predict that if such a mechanism of changing beliefs about hidden donations plays a role, the share of undisclosed offers would be higher in CHOICE-100-COSTLY compared to CHOICE-100. Moreover, donations should be lower due to the exploitation effect described above.

Table A12 shows the main outcomes of interest and compares them to the CHOICE-100 treatment. The share of undisclosed offers drops slightly, by a statistically insignificant 11 percentage points ($p = 0.248$, MWU-test), while donations, if anything increase by about 11 percent, even though this difference is also far from being significant ($p = 0.419$, MWU-test). There is no difference between the treatments for any of the relevant outcome variables. Hence, this—somewhat subtle—treatment variation does not lead to meaningful changes in market behavior. To the extent that any treatment differences could, at least in part, have been attributed to biased beliefs of buyers about hidden donations, our results of this

robustness check are in line with our interpretation when comparing CHOICE-100 with FULL INFO: Buyers in our experiment seem to hold rational beliefs when inferring donations from hidden offers.

	CHOICE-100	CHOICE-100-COSTLY	<i>p</i> -value
Share buyers who bought	0.997 (0.01)	0.996 (0.014)	0.615
Donations (offered)	44.50 (26.66)	50.51 (25.98)	0.419
Donations (sold)	43.60 (29.98)	47.90 (30.92)	0.729
% revealed	0.7465 (0.13)	0.6326 (0.26)	0.248
% revealed (sold only)	0.7561 (0.17)	0.5955 (0.32)	0.157
Prices (offered)	41.65 (12.75)	38.71 (10.75)	0.525
Prices (sold)	36.28 (13.68)	33.37 (12.68)	0.773
Payoff buyer	83.53 (13.60)	86.40 (12.93)	0.773
Payoff seller	31.25 (5.10)	29.41 (4.18)	0.356

Notes: The table reports market averages and standard deviations (in brackets) for the different treatments. The final column denotes the *p*-value from a MWU-test, testing for differences between the two treatments. The level of observation is market averages.

Table A12: Summary Statistics of Main Variables of Interest

B Experimental Instructions

— Translated from German into English —

B.1 Treatments No Info and Full Info

Welcome To Our Experiment

During the experiment, you are not allowed to use mobile phones or communicate with other participants. Please use only the computer's programs and functions intended for this experiment. Please do not talk to other participants. Should you have a question, please raise your hand. We will then come to your desk and answer your question in private. Please do not ask your question loudly. If the question is relevant to all participants, we will repeat the question and answer it. Anybody violating these rules will be excluded from the experiment and the payment.

In addition to the 4 EUR which we will pay you simply for your participation, you can earn a substantial amount of money—how much exactly depends on your decisions. We will explain this in more detail below.

Your earnings in this experiment are calculated in points. At the end of the experiment, the amount of points which you earned, will be converted to Euros. It holds that

$$\mathbf{130\ Points = 1\ Euro.}$$

General Structure

This experiment consists of 30 rounds. All of these 30 rounds are identical, which means that they follow the same rules. The points that you earn in each round will be summed up at the end for your total earnings.

In this experiment there are two different roles: buyer and seller. In addition, you will also be allocated into different groups. Each group consists of 4 sellers (Seller A, Seller B, Seller C and Seller D) and 2 buyers (Buyer X and Buyer Y). During the whole experiment, the composition of each group does not change.

At the beginning of the experiment, the computer randomly determines your role (e.g., “Buyer X” or “Seller C”) and you will be informed about it on the screen. You will keep this role for the whole experiment.

Sequence of events in the individual rounds

In each round, every buyer and every seller is given an endowment of 20 points.

Decisions of the sellers

At the beginning of the round, sellers make their decisions. Every seller makes two decisions:

1. Every seller chooses a price between 0 and 120 points at which he offers his product. In case his product is bought, the seller receives this amount of points from the buyer. For a buyer who buys the product, it has a value of 100 points.
2. Additionally, every seller decides about the production cost of the product. The production cost are between 29 and 0 points. If the seller chooses the highest production cost of 29 points, 100 points will be donated to UNICEF in case the product is sold. If the seller chooses lower production cost, the donation is reduced as well. When production cost are 0, the donation is also 0 points. The table below describes the relationship between production cost and donations:

d	100	90	80	70	60	50	40	30	20	10	0
$c(d)$	29	27	25	23	21	18	15	12	8	4	0

The **earnings of a seller** are then as follows:

- If a buyer buys the offered product, the seller earns, in addition to the endowment of 20 points, the price minus the production cost.

$$\text{Earnings if product is sold} = 20 + \text{price} - \text{production cost}$$

- If the product is not bought, the seller only earns his endowment of 20 points. No production cost accrue and no donation is generated.

$$\text{Earnings if product is **not** sold} = 20$$

It also holds that no seller can choose a combination of price and production cost which would lead to the seller making losses from selling the product.

Decisions of the buyers

After all sellers made their decisions, the buyers decide whether they want to buy one of the offered products, and if so, which one. Both sellers see on their screen the offers made

by the sellers. [NO INFO: I.e., they see the price that the seller wants for the product. The donation which this product generates is not visible to the buyers.] [FULL INFO: I.e., they see the price that the seller wants for the product and also the donation which this product generates.]

In every group one of the two buyers is chosen randomly and with equal probability. The chosen buyer then picks the one of the four products which he wants to buy, or decides to buy none of them. Then, it is the turn of the second buyer who chooses one product among the remaining ones, and this decision is then also implemented. In every round, the order is randomly determined anew.

The **earnings of a buyer** are then as follows:

- If the buyer buys one of the offered products, he earns, in addition to the endowment of 20 points, the value of the product, 100 points. From this, the price of the product is deducted.

$$\text{Earnings if product is bought} = 20 + 100 - \text{price}$$

- If the buyer does not buy one of the offered products, he earns his endowment of 20 points. No donation is generated.

$$\text{Earnings if no product is bought} = 20$$

At the end of every round, all sellers can see the offers of the other sellers and also which offers were accepted by the buyers.

The donations

The donations generated by the decisions of the participants in this experiments will be donated to UNICEF. Here, the same conversion rate of 130 Points = 1 Euro applies.

UNICEF, the United Nations Children's Fund, supports children in developing countries and areas of conflict. UNICEF promotes that children can go to school, receive medical care, clean drinking water as well as sufficient nutrition. World-wide, UNICEF takes action to protect children from exploitation and abuse.
(source: unicef.de)

The full amount from this session will be donated via bank transfer by the end of today. We will email all participants a copy of the donation certificate. If you prefer not to receive this, please let us know at the end of the experiment (e.g., when you receive your payment).

B.2 Treatments Choice-100, Choice-85 and Choice-60

Welcome To Our Experiment

During the experiment, you are not allowed to use mobile phones or communicate with other participants. Please use only the computer's programs and functions intended for this experiment. Please do not talk to other participants. Should you have a question, please raise your hand. We will then come to your desk and answer your question in private. Please do not ask your question loudly. If the question is relevant to all participants, we will repeat the question and answer it. Anybody violating these rules will be excluded from the experiment and the payment.

In addition to the 4 EUR which we will pay you simply for your participation, you can earn a substantial amount of money—how much exactly depends on your decisions. We will explain this in more detail below.

Your earnings in this experiment are calculated in points. At the end of the experiment, the amount of points which you earned, will be converted to Euros. It holds that

$$130 \text{ Points} = 1 \text{ Euro.}$$

General Structure

This experiment consists of 30 rounds. All of these 30 rounds are identical, which means that they follow the same rules. The points that you earn in each round will be summed up at the end for your total earnings.

In this experiment there are two different roles: buyer and seller. In addition, you will also be allocated into different groups. Each group consists of 4 sellers (Seller A, Seller B, Seller C and Seller D) and 2 buyers (Buyer X and Buyer Y). During the whole experiment, the composition of each group does not change.

At the beginning of the experiment, the computer randomly determines your role (e.g., "Buyer X" or "Seller C") and you will be informed about it on the screen. You will keep this role for the whole experiment.

Sequence of events in the individual rounds

In each round, every buyer and every seller is given an endowment of 20 points.

Decisions of the sellers

At the beginning of the round, sellers make their decisions. Every seller makes three decisions:

1. Every seller chooses a price between 0 and 120 points at which he offers his product. In case his product is bought, the seller receives this amount of points from the buyer. For a buyer who buys the product, it has a value of 100 points.
2. Additionally, every seller decides about the production cost of the product. The production cost are between 29 and 0 points. If the seller chooses the highest production cost of 29 points, 100 points will be donated to UNICEF in case the product is sold. If the seller chooses lower production cost, the donation is reduced as well. When production cost are 0, the donation is also 0 points. The table below describes the relationship between production cost and donations:

d	100	90	80	70	60	50	40	30	20	10	0
$c(d)$	29	27	25	23	21	18	15	12	8	4	0

3. Every seller decides whether the buyer can see the chosen donation or not. [CHOICE-100: If the seller decides to make the donation visible, it appears, together with the price, on the offer screen of the buyer. If not, the buyer cannot see which donation the product generates and will also not be told later.] [CHOICE-85 & CHOICE-60: If the seller decides to not make the donation visible, it holds that:
 - The buyer cannot see which donation the product generates and will also not be told later.

If the seller decides to make the donation visible, a random draw by the computer determines whether the donation is visible to the buyer or not:

- With a probability of [CHOICE-85:85%] [CHOICE-60: 60%] the donation is visible to the buyer and appears together with the price on the offer screen.
- With a probability of [CHOICE-85:15%] [CHOICE-60: 40%] the buyer cannot see which donation the product generates and will also not be told later.

Whether the donation is visible or not, thus depends both on the decision of the seller as well as on chance.]

The **earnings of a seller** are then as follows:

- If a buyer buys the offered product, the seller earns, in addition to the endowment of 20 points, the price minus the production cost.

$$\text{Earnings if product is sold} = 20 + \text{price} - \text{production cost}$$

- If the product is not bought, the seller only earns his endowment of 20 points. No production cost accrue and no donation is generated.

$$\text{Earnings if product is not sold} = 20$$

It also holds that no seller can choose a combination of price and production cost which would lead to the seller making losses from selling the product.

Decisions of the buyers

After all sellers made their decisions, the buyers decide whether they want to buy one of the offered products, and if so, which one. Both sellers see on their screen the offers made by the sellers. [CHOICE-100: I.e., they see the price that the seller wants for the product and also the donation which this product generates, provided the seller decided to reveal the level of the donation.] [CHOICE-85 & CHOICE-60: I.e., they see the price that the seller wants for the product and also the donation which this product generates, provided the seller decided to reveal the level of the donation and the random draw by the computer determined that the donation is visible.]

In every group one of the two buyers is chosen randomly and with equal probability. The chosen buyer then picks the one of the four products which he wants to buy, or decides to buy none of them. Then, it is the turn of the second buyer who chooses one product among the remaining ones, and this decision is then also implemented. In every round, the order is randomly determined anew.

The **earnings of a buyer** are then as follows:

- If the buyer buys one of the offered products, he earns, in addition to the endowment of 20 points, the value of the product, 100 points. From this, the price of the product is deducted.

$$\text{Earnings if product is bought} = 20 + 100 - \text{price}$$

- If the buyer does not buy one of the offered products, he earns his endowment of 20 points. No donation is generated.

$$\text{Earnings if no product is bought} = 20$$

At the end of every round, all sellers can see the offers of the other sellers and also which offers were accepted by the buyers.

The donations

The donations generated by the decisions of the participants in this experiments will be donated to UNICEF. Here, the same conversion rate of 130 Points = 1 Euro applies.

UNICEF, the United Nations Children's Fund, supports children in developing countries and areas of conflict. UNICEF promotes that children can go to school, receive medical care, clean drinking water as well as sufficient nutrition. World-wide, UNICEF takes action to protect children from exploitation and abuse.
(source: unicef.de)

The full amount from this session will be donated via bank transfer by the end of today. We will email all participants a copy of the donation certificate. If you prefer not to receive this, please let us know at the end of the experiment (e.g., when you receive your payment).

B.3 Online Replication of Full Info and Choice-100

Welcome To Our Experiment

As a first step, please close all programs and tabs that are not required for the participation in this experiment. In case you experience any problems with your internet connection during the experiment, you will have at least three minutes time to reopen the link and continue with the experiment. If you are offline for a longer time span, you will not be able to continue with the experiment. If you have any questions about the instructions, you will get the chance to ask them in our Zoom meeting after the time for reading the instructions is over. During the experiment you are not allowed to communicate with other participants.

The remaining instructions were exactly as in the laboratory experiment.