

Satisfaction Guaranteed: When Moral Hazard meets Moral Preferences*

James Andreoni[†]
University of California San Diego
& NBER

April 2017
(Current Version October 30, 2017)

Abstract

The fear of moral hazard—especially in the age of Internet commerce—can depress or prevent profitable trades. Experiments show, however, that many people prefer honesty to deceit and would not succumb to moral hazard. This paper asks a whether we can find a simple, voluntary institution that can empower moral traders, drive out amoral ones, reduce moral hazard, and restore profitable trade to markets. I find that selling goods with a Satisfaction Guarantee, accompanied by potentially minor legal or reputational enforcement, allows moral preferences to defeat moral hazard.

*I am indebted to Zachary Breig for his assistance on this project, and I am grateful to Nava Ashraf, Ernst Fehr, Oliver Hart, Caroline Hoxby, Louis Kaplow, Erzo Luttmer, Joel Sobel, Maria Titova, Matthew Weinzierl, along with editor and referees for helpful comments, and to the Google Corporation, and the National Science Foundation, grants SES1427355 and SES1658952, for financial support.

[†]University of California, San Diego, Department of Economics, 9500 Gilman Drive, La Jolla, CA 92093. (andreoni@ucsd.edu). Financial support from the National Science Foundation (SES-1427355 and SES-1658952) is gratefully acknowledged.

1 Introduction

When a buyer cannot verify the quality of a good before it is purchased—Internet transactions being a key example—moral hazard becomes a critical problem. How do buyers know they will get their money’s worth? Buyers can turn to reputational ratings, but these are often provided by the sellers, are not representative samples of buyers and, moreover, are prone to manipulation by sham raters¹. Another innovation is to allow buyers to examine the good after purchase and, if they are unsatisfied with the quality, they can return the product for a full refund. This is a practice known as *satisfaction guaranteed*.

Satisfaction guaranteed has been a common marketing strategy in the US for years, and is now also prevalent on the Internet. According to one survey, 95% of retailers have some sort of policy such that products in “like new” condition are returnable.² In addition to “like new” refunds, many retailers also accept used goods, or allow a “trial period.” In the US, many major online retailers go as far as to include a return shipping label with the merchandise—often prepaid—to lower the consumers’ transaction costs of the satisfaction guarantee. Others are more restrictive, offering a short window of time for refunds, or charging “restocking fees” upon return. Returns policies are an important component of modern marketing.³

The value of satisfaction guaranteed to sellers is easily explained by looking at two games commonly studied in the laboratory. First is the Trust game (Berg et al., 1995; Van Huyck et al., 1995) and its cousin the Gift Exchange game (Fehr et al., 1997). These games describes the fundamental moral hazard problem. Player 1, the buyer, passes money to player 2, the seller, with the hope of getting something of equal or greater value in return. Whatever is passed is scaled up (when the buyer passes in the Trust game, and when the seller passes in the Gift Exchange game), creating a surplus, which the seller can share with the buyer. The obvious equilibrium in this game is that the seller should return nothing, so should not be trusted by the buyer. As with the lemons problem (Akerlof, 1970), the market collapses.

The second game is the Ultimatum game (Güth et al., 1982). In this simple bargaining game, the proposer offers the responder a split of some surplus. If the responder accepts the offer then the division is carried out, while if the responder rejects it both sides get zero. Since something is better than nothing (by assumption), any positive offer will, in

¹See Mayzlin et al. (2014) for empirical support.

²Che (1996) summarizes a survey of Illinois retailers that report 78% give cash refunds with a receipt, 32% give cash refunds even without a receipt. Twenty-three percent limit the return period, and others limit returns to merchandise credit. However, fewer than 5% say all sales are final.

³See popular discussions about return policies at for example, *Time Magazine*, “Why a Good Return Policy Is So Important for Retailers,” Sept. 4, 2012, or *Entrepreneur Magazine*, “4 Things Shoppers Want in an Online Retailer’s Return Policy (Infographic)”, June 20, 2014.

equilibrium, be accepted. This gives the proposer all of the bargaining power; he makes the smallest possible positive offer and gains virtually all of the surplus.

Satisfaction guaranteed combines the Trust game with the Ultimatum game. Before playing the Trust game, player 2, the seller, chooses whether to offer a money back guarantee to player 1, the buyer. According to classical theory, if player 1 is offered the guarantee, knowing that she can never be worse off, player 1 will pay the “full price” to player 2. This creates a surplus which is fully in the possession of player 2. Player 2 then gives back to player 1 just enough to make it unattractive for player 1 to ask for a refund. That is, the seller gives the buyer exactly his money’s worth (or ε more), and keeps virtually all of the surplus, just like the proposer in an Ultimatum game. Satisfaction guaranteed has now (weakly) cured the moral hazard problem, making buyers slightly better off, all while allocating almost the entire surplus to the seller. If people behave according to these assumptions, then Satisfaction Guaranteed is good business and should always be adopted by sellers.

A necessary component of satisfaction guaranteed is that the promise of a refund must be credible. In order to fortify satisfaction guaranteed, in 1975 the US adopted the Magnuson-Moss Warranty Act, which specifies that representations such as “satisfaction guaranteed,” “money back guarantee,” and the like, have legal consequences. Sellers can be sued if they fail to honor them. Indeed, the webpage of the Federal Trade Commission keeps a public register of settlements with companies that have failed to comply. Nonetheless, one must ask whether it will be worthwhile for someone spending a relatively small sum on an Internet purchase, for example, to file a claim with the Federal Trade Commission. If not, then this weakens the value of satisfaction guaranteed and restores moral hazard.

Turning from theory to behavior, we know three things that should cause us to reevaluate the predictions above. In Trust games many people do actually trust sellers, and some of these sellers share the surplus equitably. However, enough sellers do succumb to the moral hazard and, on average, buyers who trust them are slightly worse off. Likewise, responders in Ultimatum games do not accept all positive offers, but typically reject “unfair” divisions. An offer of merely 20% of the surplus, for example, is very likely to be rejected, even when playing for very large stakes. Thus, moral preferences and concerns for fair play must be considered here. The existence of fair sellers means that the moral hazard problem may not be as severe as feared, and a return policy may not grant as much bargaining power to sellers as just projected.

What about enforceability? A large body of evidence shows that many people are averse to lying, or feel guilty if they have disappointed someone.⁴ Promising a satisfaction guaran-

⁴See Charness and Dufwenberg (2006), Ellingsen and Johannesson (2004), Ellingsen et al. (2010) Charness

tee, even if it is not an enforceable promise, may still be morally binding for many sellers. If a desire for honesty interacts with the offer of a guarantee, then selection into and out of a generous return policy is likely to restore some value to satisfaction guaranteed. In the end, whether satisfaction guaranteed succeeds in solving moral hazard and who benefits if it does remain open questions. The important ingredient is moral preferences; how do concerns for fair play and aversion to lying shift the bargaining power in the market?

This paper will report on a laboratory experiment that focuses on the satisfaction guaranteed game where returned items restore the pre-transaction payoffs. We find that a satisfaction guarantee that is perfectly enforced will greatly increase economic efficiency. However, sellers that share too little of the surplus are often rejected, thus undoing many efficient trades. The net effect, in contrast to predictions, is that buyers are significantly better off under satisfaction-guaranteed, but sellers' profits are about the same with and without guarantees. This is true even with experience.

We also allow sellers the option of providing their good with satisfaction guarantee. We find they overwhelmingly will do so, and those that do not are not trusted by buyers. When given the choice, therefore, sellers are far better off providing a satisfaction guarantee.

Finally, we allow fulfilling guarantees to be voluntary and non-binding. This is our most interesting treatment. We find, as expected, that nonbinding guarantees greatly reduce the trust put in sellers. This lack of trust is partly justified—of those who seek refunds, only 17% are honored. However, we also found that buyers in this condition trusted too little. Despite being strategically equivalent to the case where guarantees are not allowed, sellers in this condition who offered a guarantee were significantly more generous. In fact, they on average returned quality that was just as good as those in the condition where the guarantee is perfectly enforced. Sellers who did not offer a guarantee, by contrast, were far less trustworthy. While it was apparently under-predicted by the buyers, the selection into offering a guarantee was correlated with the trustworthiness of the seller, making it easier for buyers to have a successful exchange. Given the strict control of information in our experiment, however, there was no way other than through experience for buyers to overcome their pessimism, a constraint of the lab that could easily be overcome in reality.

What does this study teach us about how the market cures moral hazard? First we see that simple and natural institutions, such as refund policies, are highly effective in generating trust. Second, such institutions without legal constraints may nonetheless come with moral constraints that result in increased trustworthiness of sellers. Still, without some

and Dufwenberg (2010) and many others on lying and guilt aversion. Pelligra (2011) provides an interesting new psychological interpretation on this behavior, indicating empathy could be an important mediating factor between creating expectations in others and subsequently fulfilling those expectations.

legal enforcement even the moral preferences on sellers may not be enough to increase the trust of buyers. Some oversight, either from governments, courts, or market reputations may be needed to guarantee the success of “satisfaction guaranteed.” An interesting possibility is that a satisfaction guarantee could also make such enforcement through reputations much easier. In particular, that rather than building a reputation over the quality of each good a merchant provides—which could be a rather subjective and, if there are many products provided and many firms to choose from, complex task—forming a reputation for having a strong and honestly upheld return policy could be far easier for buyers to form and sellers to build. Moreover, a reputation for a good return policy could act as a signal for the overall quality of the goods and services provided.⁵

The next section will provide a brief review of the US laws on satisfaction guaranteed, and will review the relevant literature from ultimatum, trust, and gift exchange games. Section 3 presents the experimental design, and Section 4 presents the basic results. Section 5 will discuss what these results imply for contract design and enforcement. Section 6 is a conclusion.

2 Background

Here we review the econometric and experimental evidence on trust, discuss how guarantees are enforced in the US, and briefly summarize the theoretical literature on guarantees.

2.1 Trust in the Field

It has long been recognized that greater trust may enhance the efficiency of market exchange. Knack and Keefer (1997) and Zak and Knack (2001) find that countries whose residents, when surveyed, are more likely to agree that “most people can be trusted” tend to have significantly higher growth rates. Several other authors have explored similar constructs of “social capital” and made similar conclusions.

Durlauf (2002) surveys this literature and convincingly argues that inferences offered by Knack and Keefer and others may not be as evident as they suggest. He states social capital might be more productively studied with controlled experiments. Examples include the study by the anthropologist Jean Ensminger (2004) that shows a connection between trust in ultimatum games and market integration of small African villages. Barr and Serneels (2009) find positive correlations between the trust game and wages earned by workers in Ghana.

⁵The modern marketing literature has suggested this interpretation of the indirect evidence from retailers. See Janakiraman et al. (2016) for a meta-study.

A field experiment by Gneezy and List (2006) shows the positive effects of gift exchange in the labor market don't last, although later work suggests that the short duration of these effects may be asymmetric, in that efficiency *losses* after a wage *decrease* are not as fleeting (Kube et al., 2013). These studies indicate the value of institutional details that may help build trust.

2.2 Trust and Reciprocity in the Laboratory

Fehr et al. (1993) present a non-linear gift-exchange game in which “workers” have increasing marginal costs of effort and “firms” can encourage effort with efficiency wages. Positive correlations between wages and effort were observed. Berg et al. (1995), and Van Huyck et al. (1995) presented very similar models now known as the Trust game. In this linear game the proposer can pass some of his endowment to the responder, which is tripled along the way, and the responder can pass money back to the proposer at a one-for-one rate. Evidence from these games is that many people trust and many people repay that trust. However, on average trust does not pay—proposers earn back about 90% of what they passed.

What motivates people in these games? Those who repay trust must do so out of some concern for altruism and efficiency (Andreoni and Miller, 2002; Ashraf et al., 2016), aversion to inequality (Bolton and Ockenfels, 2000; Charness and Rabin, 2002; Fehr and Schmidt, 1999; Levine, 1998), an aversion to guilt (Charness and Dufwenberg, 2006, 2010) or an intrinsic taste for reciprocity (Fehr et al., 1997; Rabin, 1993).⁶ Those who exhibit trust could have two motives. First, they could care about the equity and efficiency of outcomes or, second, they could be opportunistic and take advantage of a fair or altruistic opponent (Andreoni et al., 2017; Andreoni and Samuelson, 2006). These motives must be balanced against a fear of betrayal (Bohnet and Zeckhauser, 2004).

Fehr et al. (1997), and Fehr et al. (1998) make a strong case that responders care about behaving reciprocally. Using the gift-exchange formats and proportional punishment and reward schemes they show that players respond as predicted to the behaviors of other subjects. Many have reexamined these findings, and it is a fair summary of the literature to say that negative reciprocity (punishing bad behavior) is observed consistently and often with significant effects across most studies, while positive reciprocity is relatively context dependent (see Jacobson and Petrie (2013) for a recent discussion).

⁶Sobel (2005) provides an excellent summary of the literature on trust and reciprocity. He distinguishes between two notions of reciprocity that are both central to our discussion. First is *instrumental* reciprocity, where reciprocity is intended to generate real returns in the future. This need not have any moral basis. The other notion is *intrinsic* reciprocity. This is behavior that is chosen for its own reward—reciprocating may be seen as the right or moral thing to do.

Andreoni et al. (2003b) look at environments for sharing a surplus that allow for either punishment, rewards, or both. In this linear carrot-stick environment they show that neither punishment of selfish behavior nor rewarding of selfless behavior are strong enough tools to improve cooperation, but that the two tools in combination are quite effective. This is true despite the fact that only one tool can be used at a time.

Charness and Haruvy (2002) explore preferences in a gift-exchange model and, by varying the degree of intentionality involved in offers and efforts, are able to identify that altruism, distributional concerns and reciprocity all have significant contributions to the final outcomes in these games. Cox (2004) takes a similar approach with the games of Berg et al. (1995). He builds from dictator to trust games in three steps and again finds significant roles for altruism, equity, and reciprocity. Gneezy et al. (2000) find that subjects show more trust when the potential returns are higher, indicating calculated faith in the reciprocity of others.

In all of these games, the context and costs of the reciprocal opportunities have been shown to be important. For instance, Andreoni et al. (2002) compared two sequential games with similar equilibria, but which differed in the cost of equity. They found that people tolerate inequality more when equality comes at the expense of efficiency.⁷ A different context effect is found by List (2006). He conducts a chain of studies that incrementally moves the gift-exchange game from the lab to the field. With each increment he finds behavior closer to the prediction of selfish behavior, with lower degrees of reciprocity.

Engle-Warnick and Slonim (2004) explore trust games that are repeated over time with the same partner. They find an erosion of trust when end periods are known, but less erosion when end periods are not known. Their study speaks to the importance of both reciprocity and reputations, and also to the fragile and temporal nature of trust.

Some of the most intriguing studies of trust and context relate to how social or formal enforcement of contracts can build or erode natural amounts of trust. Bohnet et al. (2001), for instance, argue that both weak and stringent enforcement of contracts achieve the greatest efficiency. Trust, they argue, is crowded out by institutions that imperfectly enforce agreements.⁸ Bohnet and Huck (2004) show that when reputations are used to build trust, the goodwill carries over to situations where reputations cannot form. Bracht and Feltovich (2008) find that allowing for voluntary enforcement (by allowing the “investor” to commit a sum to escrow in a trust game) can lead to efficient outcomes. These disparate findings are part of the motivation for this paper.

⁷This finding is evident in many studies, that is, people will prefer more for both subjects to less, even when relative allocations are uneven. Andreoni and Miller (2002) and Andreoni et al. (2003a) find significant minorities, however, are willing to “shrink” lopsided allocations toward zero for both.

⁸Gneezy and Rustichini (2000) find a related result in a field study, although here the enforcement (a fine) reduces compliance by making clear the price of non-compliance, rather than displacing trust.

Several authors have looked at the use of rating systems to build trust, with mixed results. Keser (2002) introduces a reputation management system to the standard trust game, and finds that in order to get good ratings sellers need to be more generous than splitting the surplus evenly. These effects work even if only the most recent rating is posted. Bracht and Feltovich (2009) conduct a similar study with a discrete version of the trust game, allowing buyers to observe a seller’s previous action, rather than a rating, and also allowing sellers to send cheap talk messages before the game. They find that the latter addition has no effect, while observation of actions provides a strong positive impact. Bolton et al. (2004) find that substantial improvements in efficiency from introducing online feedback mechanisms (in the form of the total number of times when the sender decided to split the surplus rather than keeping it) still do not reach the level of efficiency in the environments where parties interact repeatedly, suggesting that buyers may perceive online feedback systems as vulnerable to manipulation by the sellers. Thus, when ratings are costless and interactions can be repeated, rating improve trust. The point of the current paper is to study situations where these conditions are not met—reputations are too costly to maintain or verify, and interactions are too infrequent to benefit from incentives for repeat business.

2.3 The U.S. Laws on Satisfaction Guaranteed

The Magnuson-Moss Warranty Act of 1975 gives the US Federal Trade Commission (FTC) the authority to enforce promises of satisfaction guaranteed. It states, “A seller or manufacturer should use the terms ‘Satisfaction Guarantee,’ ‘Money Back Guarantee,’ ‘Free Trial Offer,’ or similar representations in advertising only if the seller or manufacturer, as the case may be, refunds the full purchase price of the advertised product.”⁹ Moreover, the Act makes it easier for consumers to pursue a remedy for breach of warranty in the courts, and creates a framework for resolving disputes inexpensively and informally, without litigation.¹⁰

Finding cases of successful consumer action is easy. A recent example of FTC action is the flouting of promised refunds by QT Inc., a tele-marketing company, whose promised satisfaction guarantee permits “consumers to readily obtain a full refund of the purchase price if they return the Q-Ray bracelet within 30 days.” The FTC ordered QT Inc. to turn over \$22.5 million in net profits and pay up to \$87 million in refunds to consumers.¹¹

Along a similar vein, many states in the US have enacted “lemon laws” to regulate the

⁹Magnuson-Moss Warranty Act, Title 16, Chapter 1, Subchapter B, Part 239.3, “‘Satisfaction Guarantees’ and Similar Representations in Advertising.”

¹⁰For more description of the act, go to “A Businessperson’s Guide to Federal Warranty Law” at the FTC website, <http://www.ftc.gov/bcp/online/pubs/buspubs/warranty.htm>.

¹¹This and other complaints, both large and small, are easily found on the Federal Trade Commission webpage, www.ftc.gov, and at the consumer advocate website see ConsumerAffairs.com.

sale of automobiles, both new and used, that allow buyers to request “reasonable repair attempts” after purchase.¹² Again, these laws are intended to strengthen the commitments made by sellers to ensure the quality of their products.

2.4 Economics Literature on Guarantees

A number of interesting and important papers have been written on guarantees, beginning with Heal (1977) who viewed guarantees as risk sharing arrangements. Che (1996) wrote the first theoretical paper explicitly on consumer return policies.¹³ He did not consider the moral hazard problem on the part of sellers, but assumed that consumers are uncertain about their preferences, and are risk averse. He then explored money-back guarantees as a screening method for monopoly sellers. The guarantee neutralizes risk aversion, promotes sales, and thus allows a monopolist to identify the high demand consumers *ex post*. Che shows that guarantees always improve the welfare of buyers, but monopolists offer too few of them.

Kessler and Lülfesmann (2004) consider the alternating offers bargaining model of Rubinstein (1982) with the option to return the good after purchase. In this model there is unknown quality prior to purchase and the option for multiple rounds of bargaining. Without guarantees there will be inferior quality, but equal division of the surplus. With guarantees, the moral hazard problem of the seller is solved and quality improves. However, the guarantee erodes the bargaining power of buyers and allows sellers to negotiate higher prices. The authors do not consider how moral preferences will interact with the pricing and bargaining.

3 Theory and Experimental Design

Consider a game with two players, player 1 acts as the buyer and player 2 the seller. Each player is endowed with 100 cents. We examine four conditions.

CONDITION 1: TRUST. In stage 1 player 1 passes $x \in [0, 100]$ to player 2. Player 2 receives an amount $3x$. In stage 2 player 2 observes x and can return any amount $y \in [0, 3x]$ to

¹²These vary from state to state, but a typical law stipulates what is meant by “reasonable repair attempts,” for instance that a new vehicle under warranty must be completely repaired or replaced within 18 months of being purchased. See autopedia.com for information about lemon laws across states.

¹³Papers by Mann and Wissink (1988, 1990) considered a non-strategic model of money-back guarantees, comparing them to product replacements.

A related literature on warranties also exists. See Cooper and Ross (1985) for the genesis of this literature. They view warranties as insurance policies and consider issues of double moral hazard.

In addition to game theoretic models, there is an extensive literature on money back guarantees in the marketing literature. See, for instance, Heiman et al. (2002). These papers analyze and compare the costs of various forms of refund or partial refund policies to the costs of other marketing tools, such as samples and demonstrations.

player 1. Final earnings for player 1 are $\pi_1 = 100 - x + y$, and for player 2 are $\pi_2 = 100 + 3x - y$.

While the most efficient outcome is $x = 100$, in the subgame perfect equilibrium, player 2 sets $y = 0$, hence player 1 chooses $x = 0$.¹⁴

CONDITION 2: SATISFACTION GUARANTEED. After the basic game of Trust, we now add a third “guarantee stage.” In this stage, player 1 has the option of choosing “default payoffs” rather than those earned from choices of x and y as calculated in the Trust game above. In this case the default payoffs would return *both* players to their original endowment, that is, $(\pi_1^D, \pi_2^D) = (100, 100)$.

The guarantee now alters the subgame perfect equilibrium. In the guarantee stage, player 1 would clearly choose the default if $y < x$. Hence, in stage 2 a money-maximizing player 2 chooses $y = x$, or $x + \varepsilon$. Going back to stage 1, any choice of x will yield the same payoff for player 1, that is $\pi_1 = 100$ or $100 + \varepsilon$. Hence, any amount $x \in [0, 100]$ is consistent with subgame perfect equilibrium.¹⁵ Note that with money-maximizing preferences, this multiplicity of equilibria means that a satisfaction guarantee will not assure efficiency.

What if there are moral concerns? Suppose, for instance, player 1 would prefer the default of 100 to any amount returned by player 2 that is not increasing in the amount passed, x beyond some minimal degree. It is easy to show that such an ethic, if it is common knowledge, will result in equilibria that are fully efficient. To see this, imagine that player 1 has moral preferences such that the utility of the default is not simply 100, but rather is increasing x , say $100 + \alpha(x)$ where $\alpha(0) = 0$ and $\alpha'(x) > 0$. If this is common knowledge, and as long as $\alpha < 2$, then player 1’s best response function will be to return $y = x + \alpha(x)$, which means both players 1 and 2 have payoffs that are strictly increasing in x . Anticipating this, player 1 will always choose $x = 100$. As a result, moral preferences—even if they are quite minor—are enough to reverse the prediction from the lemons problem, going from a missing market to a thriving and fully efficient one.¹⁶

¹⁴Some readers may find the payoffs $\pi_1 = 100 - x + 3y$ and $\pi_2 = 100 + x - y$ a more intuitive representation of the market transactions. While this is a defensible position, the game chosen contains the same incentives, albeit at different marginal rates, but has the clear advantage of being a game with a well-studied history.

¹⁵Note this is also a perfect equilibrium as long as the “trembles” by player 2 are independent of the amount passed by player 1. If they are increasing in the amount passed, however, then $x = 100$ could be the unique perfect equilibrium.

¹⁶Notice, the kind of moral concern just described would not be present if people applied simple outcome-based models of fairness, such as Fehr and Schmidt (1999) or Bolton and Ockenfels (2000). Models of fairness that include notions of intentions, such as Falk and Fischbacher (2006) (who employ psychological game theory as in Rabin (1993) and Dufwenberg and Kirchsteiger (2004)), however, not only capture the the intuitions in this paragraph, but also describe more accurately fairness behaviors in this and other sequentially played experimental games, as in Falk et al. (2008) and Andreoni et al. (2002). However, a simple edit to the Fehr-Schmidt model that does not require psychological game theory would result in make

The Trust and Satisfaction Guaranteed games are illustrated in Figure 1. One can easily identify the equilibria in this figure. The figure also makes salient two possible competing versions of equity. First is “equal-payoffs” in which final payoffs of the two players are the same. This should encourage player 2 to choose $y = 2x$ and encourage players to strive for the (200, 200) payoff. However, one could also justify a “split-the-surplus” notion of equity. By passing x , player 1 is creating a surplus of $3x$ for player 2, which shared evenly means $y = 1.5x$. When $x = 100$, this means a payoff of (150, 250). As we will see, both notions of equity are evident in the data.

CONDITION 3: OPTIONAL GUARANTEE. Start with Condition 2 and add a preliminary contract stage. In this stage player 2 decides whether he will provide a satisfaction guarantee. If he does, the game follows that of Condition 2 above, and if not it follows as in Condition 1. The guarantee, if chosen, is perfectly enforced.

Recall that a trustworthy seller has nothing to lose by offering a satisfaction guarantee. By contrast, an opportunistic seller may (or may not) find himself worse off in a situation with guarantees. As a result, those not offering a satisfaction guarantee will surely be mistrusted by buyers. In order to avoid revealing oneself as an opportunist, therefore, we expect all sellers to offer a satisfaction-guaranteed, and thus for this conditions to be strategically identical to Condition 2.

CONDITION 4: NONBINDING GUARANTEE. This condition adds a fifth and final stage to condition 3. In this final stage, those who offer guarantees do not have to honor them. In particular, if player 1 asks for a refund, player 2 can honor the guarantee, returning players to the (100, 100) endowment, or renege on the promise and keep the payoffs as they stand.

This last condition is the most interesting and, for many markets, the most realistic.¹⁷ Since guarantees are not enforced, this situation, without moral preferences, is strategically identical to the trust game. With moral preferences, however, people may actually trust and be trustworthy and, moreover, be averse to lying (Gneezy, 2005). If, as just discussed, market forces compel sellers to offer a satisfaction guarantee, then moral forces may compel

exactly the prediction just articulated. In particular, assume the utility the buyer gets from taking the default option would not simply be 100, but $100 + \alpha x$ where $\alpha > 0$ can be arbitrarily small (but less than 2). An ethic of this kind is rooted in the kinds of reciprocal behavior that is often observed in data that Fehr-Schmidt preferences have been evoked to explain. With this minor addition, it is trivial to derive that a Fehr-Schmidt approach would imply that the anticipated equilibrium payoff of the buyer is increasing in x .

¹⁷The design most similar to this that we know of is the “promise condition” of Glaeser et al. (2000). They gave subjects the chance to make a nonbinding promise to pass back at least what they received, that is, to promise a return ratio of 1. They found the promise had little effect, and did not improve the amount returned by player 2s.

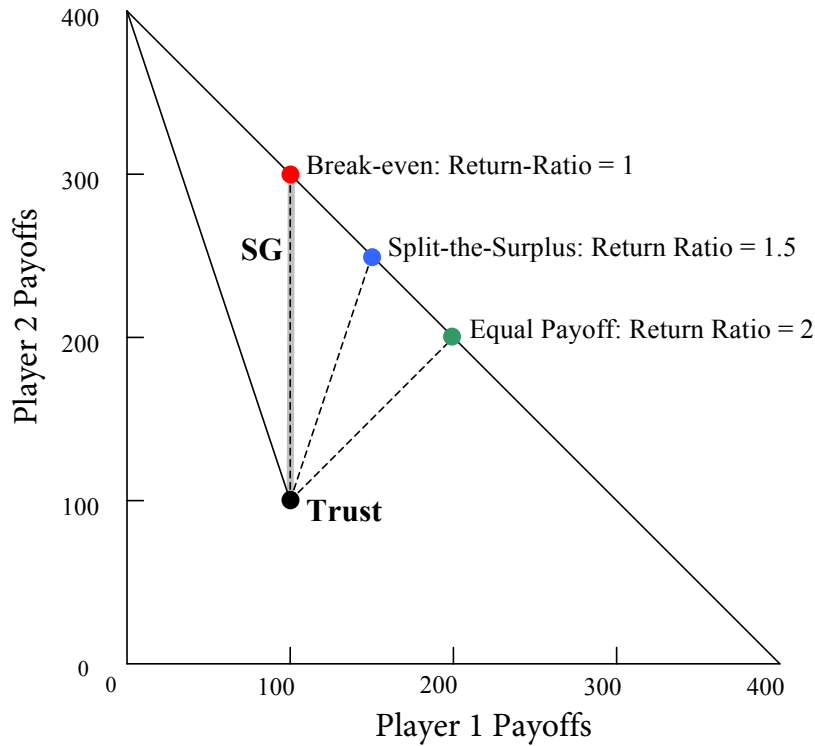


Figure 1: Subgame Perfect Equilibrium Outcomes in Trust and Satisfaction Guaranteed (SG) Games

them to honor it, in which case they should also tend to return amounts that will keep them honest and prevent a request for a refund. That is, depending on the strength of preferences of fair play and honesty, even nonbinding guarantees may increase efficiency.

3.1 The Experiment

For each session of the experiment we recruited 20 subjects. All subjects were volunteers from undergraduate economics courses. There were two sessions for each of the four conditions, meaning each condition has 40 subjects, 20 in each role, with a total of 160 subjects in the study.

Subjects interacted over a computer network. They were first presented instructions for their game (which were also read aloud to all subjects), then answered quiz questions to check their ability to calculate payoffs for both roles of buyer and seller. They were then told their own role, which they kept throughout the experiment, and began making decisions. Each session thus has 10 player 1's and 10 player 2's (called player Red and Blue in the experiment). They played 10 iterations of the game, each time with a different partner. They were told, truthfully, that they would never play the same person twice, and would

Table 1: Average Amounts Passed and Returned, by Condition

Condition	Amount Passed by Player 1		Amount Returned by Player 2	
	Rounds 1-10	Rounds 6-10	Rounds 1-10	Rounds 6-10
Trust	45	44	49	56
Satisfaction	72	80	113	130
Optional				
All	64	76	98	117
No Guarantee	15	5	4	11
Guarantee Offered	82	88	110	121
Percent Guarantees	74%	85%		
Nonbinding				
All	50	51	79	83
No Guarantee	27	18	25	4
Guarantee Offered	56	57	87	90
Percent Guarantees	81%	84%		
N	800	400	800	400

be paid for each interaction. Each subject participated in only one of the conditions above. Subjects’ instructions are included in Appendix A.

Each session lasted less than one hour. Subjects earned an average of \$15 (s.d. 4.80), ranging from \$5.13 to \$28.00. Subjects’ identities were never recorded, and all were paid anonymously and confidentially in cash at the end of the study.

4 Results

This section considers the results in light of four questions: 1) Does the satisfaction guarantee improve efficiency? 2) Who benefits? 3) Will sellers voluntarily commit to a satisfaction guarantee? and 4) If compliance is voluntary, will moral preferences (altruism, fairness, honesty and trust) be enough to sustain the efficiency properties of satisfaction guaranteed?

Table 1 presents the mean results for each condition. In what follows we will explore these data by taking advantage of the panel structure of the data to test for differences across conditions. The important comparisons will be whether Satisfaction is different from Trust, whether Optional is the same as Satisfaction, whether Nonbinding is the same as Trust, and whether offering a guarantee in the final two conditions is “good for business,” both for sellers and society.

Table 2: Amount Passed by Player 1: Two-limit Tobit Regressions with Random Effects

Independent Variable	Rounds 1-10		Rounds 6-10	
	(1)	(2)	(3)	(4)
Trust	48.010 ^{Z,S} (10.478)	47.632 ^{Z,S} (9.851)	46.967 ^{Z,S} (16.054)	46.551 ^{Z,S} (13.986)
Satisfaction	93.833 ^{Z,T} (10.677)	89.512 ^{Z,T} (10.004)	123.503 ^{Z,T} (17.382)	115.030 ^{Z,T} (15.067)
Optional	80.378 ^{Z,T} (10.593)		115.304 ^{Z,T} (16.979)	
No Guarantee		-6.147 ^{T,S} (11.417)		-21.586 ^{T,S} (20.435)
Guarantee Offered		112.687 ^{Z,T} (10.470)		135.436 ^{Z,T} (15.475)
Nonbinding	50.037 ^{Z,S} (10.490)		46.866 ^{Z,S} (16.023)	
No Guarantee		17.710 ^{t,S} (11.642)		-1.490 ^{t,S} (16.826)
Guarantee Offered		57.556 ^{Z,s} (9.962)		55.490 ^{Z,S} (14.044)
Log Likelihood	-2573.672	-2438.470	-1071.514	-995.798
N	800	800	400	400

Notes: Estimates are from two-limit Tobit regressions with random effects and Amount Passed by Player 1 as the LHS variable. Standard errors are reported in parentheses.

z & Z - Significantly different from 0 at less than 5% or 1%, respectively

t & T - Significantly different from Trust at less than 5% or 1%, respectively

s & S - Significantly different from Satisfaction at less than 5% or 1%, respectively

4.1 Does Satisfaction Guaranteed Improve Efficiency?

Table 2 presents analysis of the amount passed by player 1. Since the amount passed is bounded by 0 below and 100 above, we utilize a two-limit Tobit regressions with random effects (Wooldridge, 2010). Specification (1) simply looks at the main effects of the conditions. All the coefficients are significantly greater than zero, and the amount passed is significantly higher in conditions in which a binding guarantee is available (Satisfaction and Optional) compared to those in which it is not (Trust and Nonbinding). Column (2) shows separate estimates for the the effects of the guarantee being chosen by player 2 within the Optional and Nonbinding conditions, finding that the higher coefficient in the Optional condition from specification (1) comes largely from effect of player 2 subjects choosing to offer the guarantee. In fact, coefficients outside of the range [0 to 100] indicate that when the guarantee is offered, the median buyer will pass the maximum amount possible, but without the guaran-

tee, the same buyer passes nothing at all. A similar but less dramatic pattern occurs with the coefficients on the guarantees being offered in the Nonbinding condition. Notice too that a guarantee in Optional is not significantly different from a guarantee in Satisfaction, and a guarantee in Nonbinding is not significantly different from no guarantee in the Trust condition. No guarantee in either the Optional or Nonbinding, while strategically identical to the Trust condition, generates significantly lower amounts passed.

Columns (3) and (4) replicate the analysis of columns (1) and (2), focusing only on rounds 6-10. These regressions lead to similar estimates, with the effects of guarantees being offered in the Optional and Nonbinding conditions being even stronger.¹⁸

In sum, Table 2 indicates a strong increase in efficiency from satisfaction guarantee when it is fully enforced, and no significant effect on efficiency when they are unenforced.

4.2 Who Benefits from Satisfaction Guaranteed?

How does satisfaction guaranteed affect the distribution of payoffs among buyers and sellers? Table 3 considers *return ratios* of player 2s, that is, the amount returned divided by the amount passed, y/x , given that $x > 0$. A return ratio of 1 means player 1 breaks even, and greater than 1 yields a profit. Since return ratios are bounded between 0 and 3, we again use a two-limit Tobit to evaluate each condition, and restrict the sample to those instances in which a strictly positive amount was passed by the buyer. Column (1) controls only for the condition each of the subjects was in, and column (3) does the same but only for rounds 6-10. The coefficients on each of the conditions is significantly greater than zero.

A concern with columns (1) and (3) is that the random effect may not be independent of the choice to offer a guarantee in the Optional and Nonbinding conditions. In particular, player 2s can see the voluntary guarantee as a signal. To address this we conduct a correlated random effects analysis. Wooldridge (2010, p. 708–710) suggests that adding the mean values of the interaction of Optional and Guarantee Offered, (Mean of Opt \times Guar), or mean of the interaction of Nonbinding and Guarantee offered (Mean of Nonbind \times Guar) as regressors in the Tobit specification will correct for the correlation between the random effects and the choice to offer the guarantee.

The results are shown in columns (2) and (4) of Table 3. Here the interpretation of the coefficients on Trust and Satisfaction, and their qualitative values, are the same as in (1) and (3). The coefficients on the whether the guarantee is offered in both the Optional and Nonbinding conditions can be interpreted as the expected mean amount returned (prior to

¹⁸It is possible that subjects' experiences in previous rounds may affect how they pass. Indeed, people who have experience with higher return ratios in prior rounds tend to pass more in later rounds. Analysis of passes controlling for these effects can be found in Appendix B.

Table 3: Return Ratios: Two-limit Tobit Regressions with Correlated Random Effects, Conditioning on Player 1 Passing More Than Zero

Independent Variable	Rounds 1-10		Rounds 6-10	
	(1)	(2)	(3)	(4)
Trust	0.700 ^Z (0.133)	0.714 ^Z (0.114)	0.666 ^Z (0.156)	0.686 ^Z (0.134)
Satisfaction	1.522 ^Z (0.131)	1.522 ^Z (0.112)	1.593 ^Z (0.150)	1.593 ^Z (0.129)
Optional	1.082 ^Z (0.133)		1.220 ^Z (0.152)	
No Guarantee		-1.376 ^Z (0.470)		-1.444 ^z (0.717)
Guarantee Offered		0.255 (0.454)		-0.899 (0.681)
Nonbinding	1.042 ^Z (0.134)		0.962 ^Z (0.155)	
No Guarantee		-0.370 (0.372)		-1.234 (0.690)
Guarantee Offered		-0.281 (0.404)		-0.200 (0.697)
Correlated Random Effects:				
Mean of Opt × Guar		1.322 ^z (0.574)		2.480 ^Z (0.762)
Mean of Nonbind × Guar		1.637 ^Z (0.453)		1.433 (0.740)
Log Likelihood	-629.129	-560.543	-269.644	-249.145
N	694	694	338	338

Notes: Estimates are from two-limit Tobit regressions with correlated random effects and Return Ratio as the LHS variable. Standard errors are reported in parentheses.

z & Z - Significantly different from 0 at less than 5% or 1%, respectively

censoring) *conditional on the seller being the type that never offers a guarantee*. Adding the coefficient on Mean of Opt × Guar, or Mean of Nonbind × Guar, to Guarantee Offered then shows estimates of how a seller who always chooses to offer the guarantee would act if she were in each of these situations. It is interesting to note that both when examining all 10 rounds and when restricting the analysis to rounds 6-10, the estimated average return ratio for a seller who always offers the guarantee in the Optional condition (calculated by adding the coefficient on Guarantee Offered to the coefficient on Mean of Opt × Guar) is very similar to the estimated average in the Satisfaction condition (1.577 versus 1.522, respectively, in rounds 1-10, and 1.581 versus 1.593 in rounds 6-10).

Table 2 showed that buyers passed significantly more when the guarantee was offered in the Nonbinding condition than when it was not, despite the two being strategically equivalent. This suggests that they believed the offer of a Guarantee contained some signal value about the intentions of the seller. Was this belief justified? In Table 3 the fact that the coefficient on Mean of Nonbind \times Guar was significantly greater than zero for rounds 1–10 suggests that it was. The fact that this coefficient becomes smaller and loses significance in later rounds is also interesting—it suggests that as those who tend to give less learn over time that the guarantee will cause buyers to pass more, they offer the guarantee more often, reducing the signal value of the guarantee.

It has been shown by others that the return ratio can be influenced by the amount sent (e.g. Glaeser et al. (2000)). In particular, the effects noted in Table 2 may simply be a result of the different passes by buyers in each of the conditions. We account for these effects in Table 4. Columns (1) and (3) are similar to the same columns in Table 3, but allow for different linear effects of passes in each condition. We see that in each condition the coefficient on the interaction with Pass is significantly different from 0 with p -values of 0.01, all with magnitudes of 0.003 to 0.006, which are not statistically distinguishable from each other. This means, for instance, that passing 100 rather than 50 will increase the return ratio by 0.2 to 0.3 in each condition. In Trust this is not enough to make the predicted return ratio greater than one when passing 100, while in all other conditions passing 100 is predicted to be profitable. As with Table 3, the coefficient on Satisfaction in Table 4 is again significantly greater than each of the other coefficients at the 5% level, while none of the other conditions are statistically different from each other.

Columns (2) and (4) of Table 4 again make the correction for correlated random effects. We see that for Optional and Nonbinding those selecting into the guarantee are primarily responsible for the higher return ratios. In column (4) the coefficients become unstable because in rounds 6–10, buyers only passed a positive amount six times when the guarantee was not offered in the Nonbinding condition.

Table 5 reports the average earnings for the two players. We see that the buyer, player 1, is far better off under Satisfaction Guaranteed. Earnings increase from 96 to 138 per round, a rise of 44%. Over the last five rounds the difference is 48%. By contrast to player 1 and to the standard prediction, the seller, player 2, is actually worse off over all 10 rounds. However, for just the last 5 rounds, average earnings by player 2 in Trust and Satisfaction are nearly identical, 189 versus 188. Total earnings in Satisfaction are higher, rising from 290 to 314 overall, and from 288 to 335 for the final 5 rounds. This is an increase of 8 to 16%. When expressed as a gains-from-trade rather than earnings (that is, subtracting 200 from the base), this is an increase in the surplus of 26 to 53%.

Table 4: Return Ratios: Two-limit Tobit Regressions with Correlated Random Effects Conditional on Player 1 Passing more than Zero

Independent Variable	Rounds 1-10		Rounds 6-10	
	(1)	(2)	(3)	(4)
Trust	0.393 ^Z	0.410 ^Z	0.335	0.351 ^z
	(0.146)	(0.129)	(0.174)	(0.155)
Pass × Trust	0.006 ^Z	0.006 ^Z	0.006 ^Z	0.006 ^Z
	(0.001)	(0.001)	(0.002)	(0.001)
Satisfaction	1.231 ^Z	1.231 ^Z	1.341 ^Z	1.340 ^Z
	(0.155)	(0.137)	(0.187)	(0.170)
Pass × Satisfaction	0.004 ^Z	0.004 ^Z	0.003 ^z	0.003 ^z
	(0.001)	(0.001)	(0.002)	(0.001)
Optional	0.609 ^Z		1.002 ^Z	
	(0.163)		(0.205)	
Pass × Optional	0.006 ^Z		0.003	
	(0.001)		(0.002)	
No Guarantee		-1.191 ^z		-51.870
		(0.517)		(1681.049)
Pass × No Guarantee		-0.005		1.495
		(0.007)		(49.443)
Guarantee Offered		0.068		-0.957
		(0.465)		(0.682)
Pass × Guarantee Offered		0.002		0.001
		(0.001)		(0.002)
Nonbinding	0.732 ^Z		0.495 ^Z	
	(0.156)		(0.189)	
Pass × Nonbinding	0.005 ^Z		0.007 ^Z	
	(0.001)		(0.002)	
No Guarantee		-0.797		-3.373 ^Z
		(0.430)		(1.221)
Pass × No Guarantee		0.008 ^z		0.032 ^z
		(0.004)		(0.013)
Guarantee Offered		-0.714		-1.044
		(0.417)		(0.679)
Pass × Guarantee Offered		0.004 ^Z		0.005 ^Z
		(0.001)		(0.002)
Correlated Random Effects:				
Mean of Opt × Guar		1.341 ^z		2.438 ^Z
		(0.570)		(0.743)
Mean of Nonbind × Guar		1.785 ^Z		1.947 ^Z
		(0.457)		(0.708)
Log Likelihood	-593.882	-531.798	-250.979	-227.843
N	694	694	338	338

Notes: Estimates are from two-limit Tobit regressions with correlated random effects and Return Ratio as the LHS variable. Standard errors are reported in parentheses.

z & Z - Significantly different from 0 at less than 5% or 1%, respectively

Table 5: Average Earnings For Player 1 and Player 2, by Condition.

Condition	Rounds 1-10			Rounds 6-10		
	Player 1	Player 2	Total	Player 1	Player 2	Total
Trust	96	194	290	99	189	288
Satisfaction	138	176	314	147	188	335
Optional						
All	119	174	293	130	182	312
No Guarantee	86	144	231	97	112	210
Guarantee Offered	131	184	315	135	194	329
Nonbinding						
All	112	186	298	110	189	299
No Guarantee	88	166	252	83	154	237
Guarantee Offered	118	191	308	115	196	311


The differences between Trust and Satisfaction Guaranteed can be seen in Figures 2. This shows the frequencies of outcomes over the final five rounds. Each circle is centered on a point in the data, and the larger the circle the more observations at that point. Figure 2a shows significant misplaced trust in the Trust condition, and many instances of disappointed player 1s. Figure 2b shows the clear improvement from satisfaction guaranteed. The guarantee ensures that in none of the observations is player 1 worse off than at the endowment point, and large numbers of interactions resulted in equitable outcomes of equal-payoffs (200, 200) and split-the-surplus (150, 250).

What does this say about the institution of satisfaction guaranteed? Focusing on the final 5 rounds, this indicates a big gain for player 1, the buyer, no net impact on player 2, the seller, and a significant 53% increase in realized gains from trade.


4.3 Do Sellers Voluntarily Commit to Satisfaction Guaranteed?

What happens when we allow subjects themselves to determine whether they will offer a contract with a satisfaction guarantee? We predict, in light of the results above, that all subjects should offer the guarantee. Table 6 shows that over all rounds, subjects in the Optional condition offer guarantees 74% of the time, rising to 85% by the final 5 periods. Nineteen subjects offer guarantees at least 5 of the 10 rounds. Although only 4 of the 20 player 2s offered the guarantee all 10 rounds, 11 subjects offered them in all of the last 5 rounds. In the final round 18 of 20 subjects gave the guarantee.

This is evidence that subjects are learning they are better off offering a guarantee than not. Returning to Table 2, we see that in the Optional condition the amount passed is significantly higher with a guarantee, and that this difference increases by the latter half of



(a) Trust, Last Five Rounds



(b) Satisfaction Guaranteed, Last Five Rounds

Figure 2: Increase in Efficiency from Trust to Satisfaction Guaranteed

Table 6: Percent of Player 2s Offering Satisfaction Guaranteed.*

Condition	Percent Who Offered Guarantee	Percent of Guaranteed to Request Refund	Percent of Refund Requests Honored
<i>Rounds 1-10</i>			
Satisfaction	100%**	23%	100%*
Optional	74%	28%	100%*
Nonbinding	81%	25%	17%
<i>Rounds 6-10</i>			
Satisfaction	100%*	18%	100%*
Optional	85%	26%	100%*
Nonbinding	84%	26%	14%

* 200 observations by condition for rounds 1-10.

** 100% is by experimental design.

the game. But does the offer of a guarantee really matter to the returns? Table 3 shows that those offering guarantees more often average significantly higher return ratio than those who tend not to offer guarantees. Thus, those offering guarantees are both treated better by buyers, and respond more generously as sellers. Table 5 shows that sellers (player 2s) who offer guarantees make almost 30% more than those who don't. Over the last 5 rounds the gap is almost 75%. Both of these differences are significant.¹⁹

It is interesting to compare the Optional condition to the Satisfaction Guaranteed condition. Return to Table 5 and compare the earnings for Satisfaction under to the earnings for Optional under Guarantee Offered. These numbers are very similar for both players 1 and 2. This is a curious juxtaposition with the finding discussed in the prior paragraph. The fact that not all people are offering guarantees might suggest that the “cheats” are revealing themselves, leaving a population of more trustworthy people among those who offer guarantees. This appears to not be precisely true—those who offer the guarantee are not more trustworthy than their counterparts in Satisfaction, but those who do not offer the guarantee are less trustworthy.

The similarities between Satisfaction and Optional conditions can be seen by comparing Figure 3 below with Figure 2b above. These both show the last five rounds of play. The similarity in the patterns is striking.²⁰ Removing those interactions in which the guarantee was not offered, marked by the shaded circles in Figure 3 makes the comparison even more precise. This will be a difference with the next game.

¹⁹For all rounds $z = 9.51$, and for the last 5 rounds $z = 9.19$.

²⁰Both the amounts passed and the return ratios can be shown to be not significantly different between these two. However, joint tests find significant differences at standard ($p \leq 0.05$) significant levels.

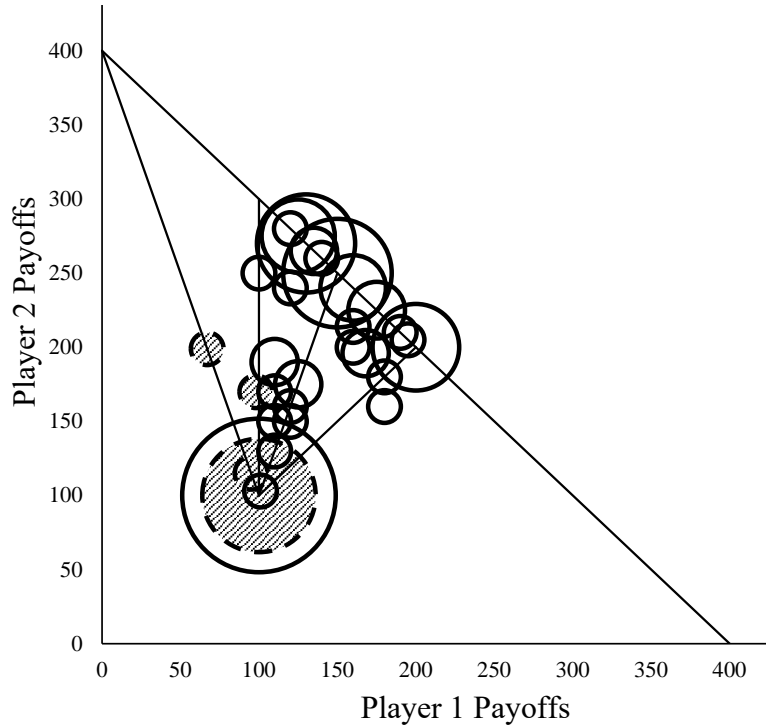


Figure 3: Optional Guarantee, Last Five Rounds. Shaded areas indicate the guarantee was not offered.

4.4 Caveat Emptor: Will Nonbinding Guarantees Still Improve Efficiency?

We now consider the Nonbinding game, which is the most complex and interesting version of the satisfaction guaranteed game. Here if the seller chooses to offer a guarantee, and the buyer requests a refund, the seller can renege on the promise and deny the refund.

Begin with the preliminary contract stage. As with Optional, most players offer the guarantee, with two main differences. First, when the default is not binding, sellers offer it much more freely. Nine of 20 subjects offered the default all 10 rounds—more than twice the rate for Optional—and 16 offered it 8 rounds or more. A second difference is on the opposite extreme. Two of the subjects chose *never* to offer the default. We asked subjects in the post-experiment questionnaire to explain their actions, but neither subject gave any insight into this decision.²¹ When we look ahead to how these two behaved when they were passed positive amounts, we get a clue. Between the two of them they were offered

²¹One subject said nothing, and the other said simply, “Never gave the default option,” which was our language for the satisfaction guarantee. It is doubtful that these two subjects did not understand the instructions. Quizzes given before each session required subjects to correctly calculate payoffs in three examples before moving on to the game (see Appendix A). We are left, therefore, to speculate about their motives.

positive amounts 11 times (an average positive offer of 54), but only returned a positive amount 1 time (returned 20 when passed 30 in round 3). It seems, therefore, that neither had intentions of returning anything they received. Hence, it is possible that these two were “honest thieves”—they did not want to tell a lie by promising a guarantee that they would not honor (Gneezy et al., 2013; Hurkens and Kartik, 2009).

Next look at the actions of player 1. In Table 2 we see the amount passed in the Nonbinding condition is significantly lower than in the Optional and Satisfaction conditions, but not significantly different from the Trust condition. This is true even when the guarantee was offered.²² Player 1’s actions in Nonbinding are not significantly different from player 1’s actions in Trust—buyers place no extra trust in a nonbinding guarantee.

Turning to player 2, we ask whether player 1s should have placed more trust in player 2s. Here there is evidence that they should have. Looking at Table 3 the actions of player 2s appear to be between those in Trust and Satisfaction. In fact, average return ratios for the Nonbinding condition are above 1 when looking at all 10 rounds, indicating a profit opportunity for buyers. Looking at return ratios under guarantees, those in the Nonbinding condition are not significantly different from those in Trust or Optional, but are significantly less than those in Satisfaction.²³

This can be clarified by an example. Suppose, player 1 sends 100 if the guarantee is offered (0 otherwise), and requests refunds if the returns are unprofitable. Given the average return rates observed, 82 percent of his offers will be profitable, earning an expected 147, 3.1 percent will earn refunds (17 percent of requests), earning 100, and 14.8 percent will not get refunds, yielding just 32. Altogether, fully trusting an unenforced guarantee yields an expected payment of 128.

The actions of player 2s now contrast strikingly with those of player 1s. Player 2s are much more trustworthy than their counterparts in the Trust condition. Even though they are not as trustworthy as those in the Satisfaction condition, they are trustworthy enough that profitable exchanges are possible on average.

Perhaps player 1 offered less because of a fear of variance, that is, risk aversion.²⁴ Table 6 shows the fraction of player 2s who offer guarantees, the fraction of those offers that generate a refund request, and the percent of those requests that are honored. The first column shows

²²A test for equality of Nonbinding and Satisfaction (Nonbinding and Optional) for all rounds has a p -value of 0.0034 (0.0418), and for the last five rounds has p -value 0.0012 (0.0035). Comparing Satisfaction (Optional with guarantees) to Nonbinding with guarantees, the p -value is 0.0235 (0.0001) in rounds 1-10, and 0.0039 (0.0001) when restricting attention to rounds 6-10.

²³The test of equality between Satisfaction and Nonbinding has a p -value of 0.0103 in rounds 1-10, and 0.0034 in rounds 6-10.

²⁴Eckel and Wilson (2004) demonstrate that there is a weak inverse relationship between trust and risk aversion.




Figure 4: Nonbinding Guarantee, Last Five Rounds. Shaded areas indicate the guarantee was not offered.

that Optional and Nonbinding conditions are fairly similar, and the second column shows the frequency of refund requests is also quite similar across Satisfaction, Optional, and Nonbinding. However, the third column shows a huge difference. Of the 40 requests for refunds in Nonbinding, only 7 were honored. For the last 5 rounds, only 3 of 21 requests actually received a refund. Looking within subjects, the only subjects who seemed unambiguously honest in their offers of guarantees were the two subjects who never offered them. Subjects who got more than one request for a refund all denied at least one of them.²⁵

What is the net effect on earnings? Table 5 shows that player 1s, the buyers, do far worse in the Nonbinding condition than in either Satisfaction or Optional. This is even true when conditioning on the presence of a guarantee. By the last 5 rounds the difference in earnings between the Trust and Nonbinding conditions is insignificant for player 1s. Looking at player 2s, the sellers, their payoff is nearly exactly the same on average, regardless of the condition. In sum, the introduction of nonbinding guarantees does little to improve efficiency—overall the improvement is not statistically significant.

The result can again be seen graphically. Figure 4 shows the pattern of outcomes for the final five rounds of the Nonbinding condition. Note the gray circles distinguish cases

²⁵Only one subject honored all requests, but it's a trivial case. This subject got a single request. The amount passed was 3 and returned was 4, so only 5 cents was lost by player 2.

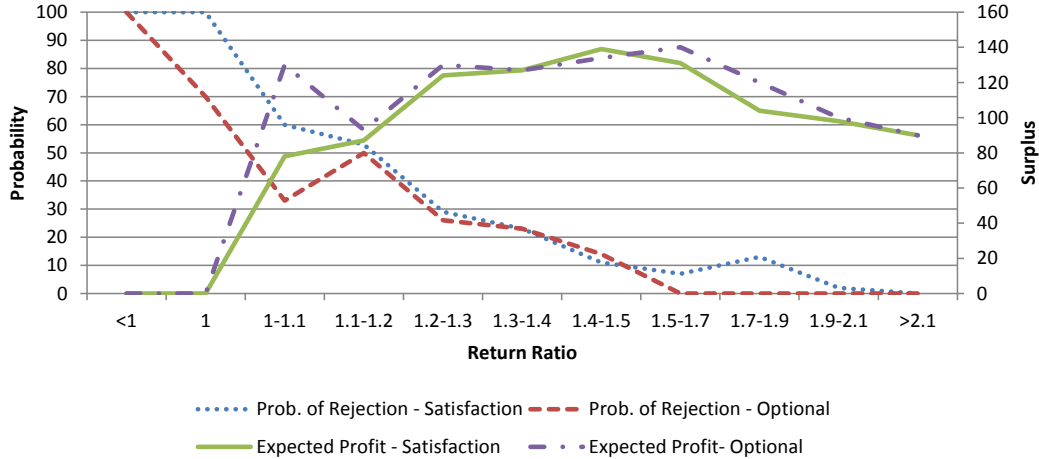


Figure 5: Probability of Requesting a Refund, and Seller’s Expected Surplus, Conditional on Return Ratios in the Satisfaction Guarantee Condition

where no guarantee was not offered. While containing some of the shades seen in Figure 2b from Satisfaction, it most resembles the outcomes from Trust seen in Figure 2a. Note the contrast of this with the surprising results of Bohnet et al. (2001) and Fehr and List (2004), who show that zero enforcement can be more efficient than imperfect enforcement. Here, satisfaction guaranteed with no enforcement provides no improvement in efficiency over no guarantee at all.

5 Trust, Reciprocity and the Law

In this section we address two issues about the interactions of trust and reciprocity with enforcement. First we consider how much fairness and reciprocity are driving the efficiency of the satisfaction guarantee. These notions have figured prominently in the work of Fehr et al. (1997), Fehr et al. (2007), and Brown et al. (2004) for instance, who state that fairness and reciprocity are potent enforcement devices.

Second we step back and look at all four institutions above at once and get a more complete picture of how satisfaction guarantees are altering the behavior and expectations of both the buyers and sellers.

5.1 Fairness and Reciprocity in Satisfaction Guarantee

In Section 3 we made the theoretical point that a satisfaction guarantee will assure efficiency if buyers will reject trades that, while profitable, do not give a sufficiently fair return. Figure 5 (left axis) shows the probability of requesting a refund in Satisfaction Guarantee condition,

conditional on the return ratio. Letting r be the return ratio, then we see, as expected, all unprofitable return ratios, $r < 1$, result in refunds, as do all “break even” return ratios, $r = 1$, when the condition is Satisfaction. However, many profitable return ratios, $r > 1$, also result in a refund. A seller who chooses a return ratio of 1.2, for instance, will have a greater than 50% chance of having to give a refund in both the Satisfaction condition and when the guarantee is offered in the Optional condition. If the buyer passed all 100 to the seller, such a refund means forfeiting net gains of 180 for the seller and 20 for the buyer. In the Satisfaction condition, the probability of a refund stays positive until return ratios exceed 2. So when the guarantees are enforceable, moral preferences are playing an important role in driving their use toward efficiency.

If we think of fairness of buyers as a constraint on sellers, then we can ask, what return ratio should a money-maximizing seller adopt? As shown in Figure 5 (right axis), the most profitable return ratio is between 1.4 and 1.7, depending on the condition, a range including the split-the-surplus ratio of 1.5.²⁶ Even so, these sellers should expect about 10% of customers at this return ratio to seek a refund. Notice that a supplier who is choosing the profit maximizing r will average earnings of about 240. This far exceeds the average earnings in the Trust condition of 194. Given the choice, therefore, adopting a binding institution of Satisfaction Guarantees seems far superior for sellers than not.

As with previous research, this confirms that tastes for fair play—when cleverly combined with a simple marketing innovation—are indeed driving efficiency in the market.

5.2 Voluntary Contracts and Voluntary Compliance

Our four treatments can be paired into two groups that, in the absence of moral preferences, are virtually equivalent strategically. First are Satisfaction Guarantee and the Optional Guarantee, and second are Nonbinding and Trust conditions. In this section we explore when and how moral preferences might break these similarities.

Look first at the amounts passed. Figure 6 shows the distribution of the amount passed, given that guarantees were offered in the Optional and Nonbinding treatments. Here we see evidence of the expected equivalence. The first two bars over each category compare the Satisfaction Guarantee with the Optional Guarantee. These two are quite similar. If anything, those in the Optional condition are more generous than those in which the guarantee is required. Depending on the test used, the difference between these two conditions is either

²⁶Appendix C confirms this result using regression analysis to solve for optimal return ratios in the two conditions for which the guarantee was binding. The estimated optimal return ratios fall between 1.55 and 1.65.

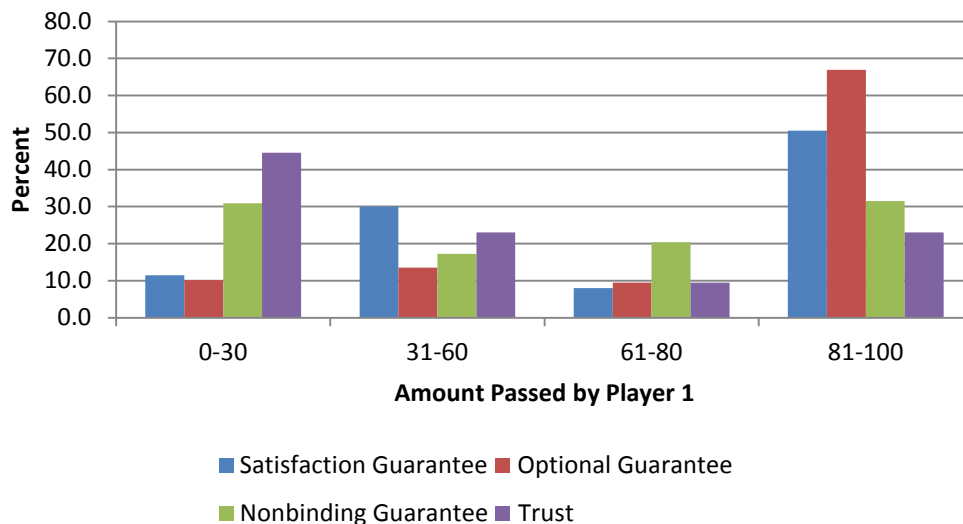


Figure 6: Distribution of Amount Passed by Player 1, Conditional on Guarantees Offered in Optional and Nonbinding conditions

not significant or marginally significant.²⁷

The final two bars over each category in Figure 6 compare Nonbinding Guarantee and Trust. Again we see the predicted similarity—the two are not significantly different by any of the tests used.²⁸ Across the two sets, however, Trust and Nonbinding treatments are different from both Satisfaction and Optional treatments.²⁹ This reinforces what was found in Table 2; that player 1 is treating the nonbinding guarantees as meaningless.

Are the refunds offered in Nonbinding actually meaningless? Figure 7 shows the probability of requesting a refund conditional on the return ratio offered. We see again that Satisfaction and Optional Guarantee are very similar. The difference between them is not significant. The Nonbinding condition is, by contrast, well below the other two, as buyers appear to have lower expectations.³⁰ The last line in Figure 7 shows how often a refund request is honored in Nonbinding. As can be seen, the promised refund is not often honored.

²⁷We test this in two ways, which we use in all the footnotes to follow in this section. First, we organize the data by subject, finding the average amount passed for each. We then compare the distributions of subjects' average choices using Mann-Whitney tests ($z = 1.89$), and Kolmogorov-Smirnov tests (p -value 0.06).

²⁸Organizing the data by subject, as in the previous footnote, and using Mann-Whitney tests ($z = 0.87$) or Kolmogorov-Smirnov tests (p -value 0.135) we see that in each case the difference in distributions is not significant.

²⁹For Trust versus Satisfaction, Kolmogorov-Smirnov (KS: p -value 0.023) and Mann-Whitney (MW: $z = 2.88$) both indicate significance. For Trust versus Optional the differences are significant as well (KS: p -value 0.003, MW: $z = 3.6$). For Nonbinding versus Optional the difference is significant (KS: p -value 0.008, and MW: $z = 2.89$), but for Nonbinding versus Satisfaction, the difference is insignificant or marginally significant (KS: p -value 0.275, MW: $z = 1.72$).

³⁰A return ratio of 1, for instance, is certain to get a refund request in Satisfaction, but only faces a 28% chance of a request in Nonbinding. Surprisingly, even unprofitable return ratios have only an 80% chance of generating a refund request in Nonbinding.

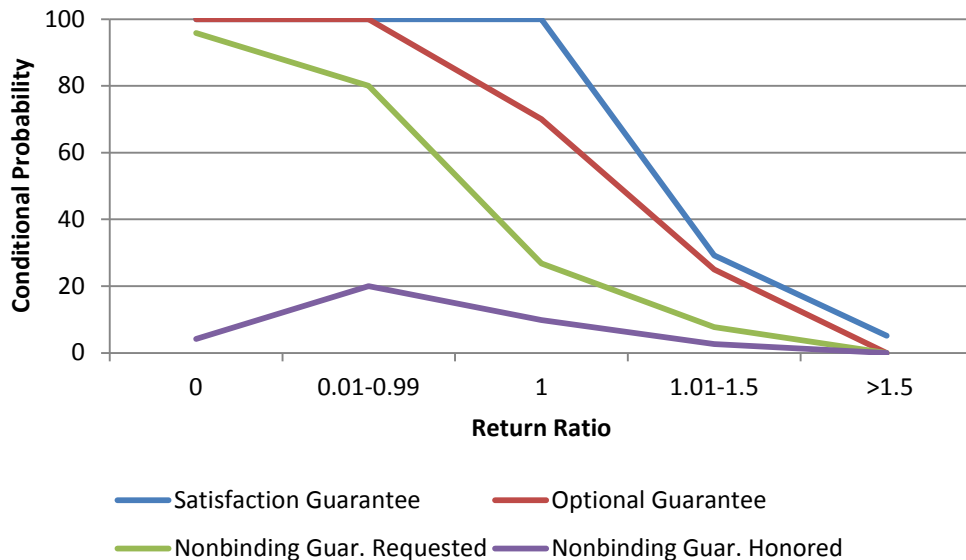


Figure 7: Conditional Probability of Requesting a Refund

So, indeed, the promise is almost meaningless—at least to those who are treated poorly enough to request a refund. Is it still possible that the promise has value? That is, do moral preferences lead those promising refunds to be so generous as to make asking for one unnecessary?³¹ Consider Figure 8. Here we look at the distribution of return ratios across treatments, again under the condition that guarantees are offered in Optional and Nonbinding treatments. As above, we should expect Satisfaction and Optional to be similar, and one can clearly see that they are.³² Comparing Nonbinding and Trust, however, the comparison is not as expected. The two are not similar, and the difference is highly significant.³³ By contrast, the statistical tests comparing the Nonbinding to the Optional treatments—which are predicted to be different—show the two are nearly indistinguishable.³⁴ This means that the sellers in the Nonbinding treatment are behaving nearly identically to sellers in the Optional treatment for whom the guarantee is binding. It appears that those who plan to return less are less willing to offer a guarantee. One hypothesis is that sellers do not want to face their own deceit.³⁵ Whatever the reason, we confirm the findings of Tables 3 and 4 that offering

³¹This contrasts with the “promise condition” of Glaeser et al. (2000). Sellers only promised to send back at least what they received, and no refunds were possible. This promise, however, did not generate extra generosity.

³²While they appear similar in the figure, the difference is marginally significant by a Kolmogorov-Smirnov test (p -value .059) and significant in a Mann-Whitney test ($z = 2.07$).

³³A Kolmogorov-Smirnov test (p -value 0.017) and Mann-Whitney test ($z = 2.24$) both indicate a significant difference.

³⁴Kolmogorov-Smirnov (p -value 0.56) and Mann-Whitney ($z = 0.22$) test indicate only minor differences between the two.

³⁵In principle this hypothesis can be tested by including a condition in which it is mandatory to offer a satisfaction guarantee, but optional to honor it. This would presumably remove the moral constraint of

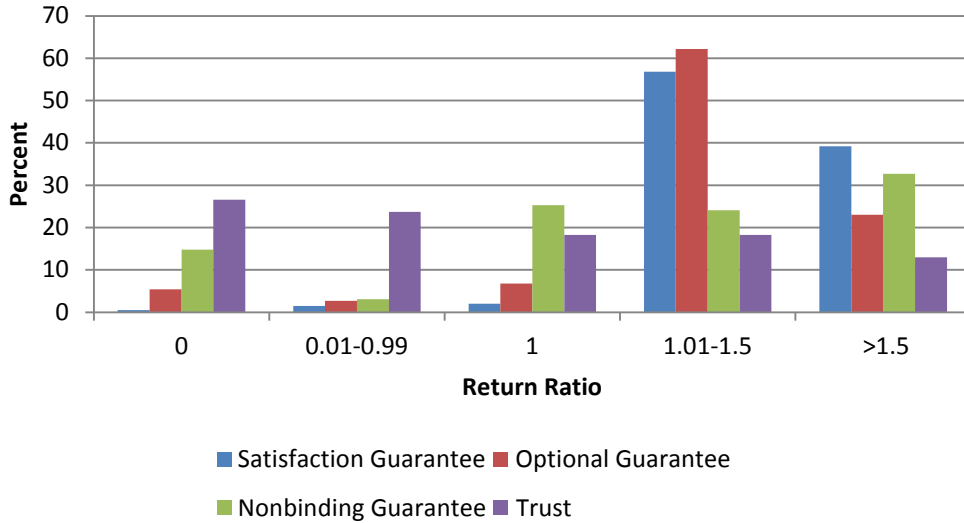


Figure 8: Distribution of Return Ratios Across Conditions

a nonbinding guarantee is correlated with more generous return ratios.

In sum, there is a minor paradox. Those promising a guarantee that they are not required to honor nonetheless act in a way that is statistically indistinguishable from those for whom the promise is binding. Even so, buyers have insufficient faith in the moral requirements put on sellers by this promise. As a result, a satisfaction guarantee without any binding enforcement fails to increase efficiency. The resolution of this paradox may rest in the abstractions of laboratory experimentation. Here we impose the problems of the market we wish to study: costly moral hazard, no opportunities for reputations, and no way for the truly trustworthy sellers to distinguish themselves. The experiments suggest that a modicum of extra trust in non-binding guarantees could generate significant gains in efficiency. This could accumulate with more experience than we are able to offer in the lab, or from return policies themselves gaining reputations as signals of quality.³⁶ If honest sellers do indeed want to use lenient return policies as signals of quality, then the burden could fall on them to monitor and expose fraudulent sellers who would undermine their signal.

6 Conclusion

This paper was designed to look at a realistic market innovation to promote trust, trade, and efficiency. We ask whether and how a satisfaction guarantee can improve economic efficiency.

having told a lie. Exploring this and other aspects of guarantees is left to future work.

³⁶Informal evidence from retailers supports this conjecture. See the National Retail Federation’s column, “Happy Returns,” such as <https://nrf.com/news/technology/happy-returns>, where they offer the advice that, “retailers with more lenient return policies fare better overall.”

Are moral preferences enough to defeat moral hazard, or does the policy need legal backing?

The experiment explored satisfaction guarantees in three stages. First, they are mandatory and enforced—all sellers must offer and honor them. Second, they are optional, but fully enforced. Third, they are optional but unenforced, a *caveat emptor*.

We find four key results. First, when enforced, satisfaction guarantees can dramatically increase efficiency and reduce moral hazard. Opposite to predictions from the model without moral preferences, however, sellers are no worse off by offering the guarantees, but buyers benefit greatly. This suggests that the guarantee is interacting with preferences for fairness and equity in important ways—a fact we validate empirically.

Second, when guarantees are optional sellers that don't offer them are not trusted nearly as much as sellers that do. However, the choice of offering a guarantee was also revealing in our data; those not offering guarantees were much less trustworthy, and less trusted.

Third, we find that nonbinding guarantees invite a number of effects that, in their own ways, could be interpreted as expressions of heterogeneous moral preference. One of these effects is that, we get the appearance of subjects who have a clear intention to succumb to moral hazard. Interestingly, many of these sellers also do not offer the nonbinding guarantee. As such they are rarely trusted, making their behavior a kin to exiting the market. By selecting out of the market, these “honest thieves,” leave the pool of sellers more favorable to consumers. Of those who offer the nonbinding guarantee, there are two opposite effects. On one side is the emergence of a small minority of opportunists. On the other side, we see positive effects of agency on the behavior of the honest sellers. Because they must face both moral choices of offering and honoring a guarantee, compared to those in the other two conditions with guarantees, these sellers are more generous to buyers and as a result get fewer requests for refunds. When refunds are requested, especially for return ratios above 1, however, these sellers often refuse to honor them, perhaps because of a moral determination that they indeed treated the buyer fairly. Adding these effects together, the average buyer can expect to be better off by trusting the sellers even with non-binding guarantees. Unlike when guarantees are enforced, these sellers face a real risk of being fleeced, but also the more likely outcome of being well-rewarded. On average, moral preferences have counteracted moral hazard among sellers.

Fourth, despite the discussion in the prior paragraph, when the guarantee is not binding, buyers don't trust sellers enough. The selection into offering a guarantee has two effects. First, more sellers adopt very favorable return ratios than when guarantees are required. It appears the promise of a satisfaction guarantee evokes stronger moral responsibilities under self-selection, and this works to improve efficiency. The second thing it does is it allows “honest thieves” to self-identify by choosing never to offer a guarantee and to simultaneously

(and perhaps morally justifiably) adopt very low return rates. Again, those who intend on acting selfishly but prefer honesty to deceit have moral preferences that limit the social costs of their concession to moral hazard. Thus for both the generous and selfish types of sellers, moral preferences seem to be counteracting moral hazard.

What have we learned from this? First, we have learned that moral hazard is easily overcome with the simple market innovation of a satisfaction guarantee. When this is perfectly enforced, efficiency improves dramatically. When it is selected endogenously, sellers quickly understand its value. When it is both voluntarily adopted and honored, the net results are less clear. On the one hand, moral preferences help make the landscape more favorable for buyers by curbing moral hazard of those sellers who would behave selfishly, and by enhancing generosity of those sellers who find pleasure or pride in keeping promises. Yet, within the scope of our experiment, buyers' collective degree of trust did not also improve with to account for the moral preferences of sellers. Given the data, we cannot conclude that non-binding guarantees are a success, but neither are they a failure. Perhaps with more time to experiment and gain experience, or with the opportunity to share even small bits of information either privately or publicly, the buyers could discover that they can profit by trusting more.

In addition to experience, other pressures would also seem likely to help buyers discover the opportunity for benefit from non-binding return policies. Foremost are reputations and selection. Businesses that routinely flout their guarantees may, at the very least, lose repeat business. Likewise, firms known for honoring guarantees may attract customers. However, if guarantees can be enforced by reputations, then it suggests that quality can be enforced this way as well. However, "quality" can be personal and subjective. Moreover, firms often have many products, and these products are often changing over time. Both of these make reputations more difficult to form and hold over the actual goods or services sold. However, the leniency of a return policy and the buyer's experience with similar return policies from other merchants could allow the buyer a low cost method of identifying return policies that are correlated with trustworthiness. Even if shoppers are dealing with new sellers for every interaction, basing the trust of the seller on the "reputation" of the return policy rather than on reports on the quality of the goods sold may be a relative easy reputation for both the buyer and seller to maintain. If this is true, then it has another advantage. Firms that offer similar return policies and do so honestly will have incentive to protect the integrity of their return policy by exposing amoral sellers who abuse the guarantee, or perhaps even reporting malfeasance to the FTC. Thus, reputations based on offers of satisfaction guaranteed could have three nice properties: every seller should offer one, return ratios should always be profitable for buyers, and the industry has an incentive to self-police competitors who falsely

promise guarantees.

A third observation is that, while satisfaction guarantees predominate US markets, they are less common in other countries. It could be that the US has stronger enforcement through the Magnuson-Moss Act, and so achieved a different equilibrium. Alternatively, there could be different constellations of preferences for equity or tolerance for opportunism in different parts of the world that affect the degree of moral hazard in the first place. As e-commerce grows, however, return policies are becoming more common around the globe.

This paper also raises the prospect of considering the broader panoply of satisfaction guarantees. For instance, many firms sell goods with a “free trial offer” or promising “double your money back.” How would this affect the bargain between buyer and seller, and the signal sent about quality? How can such offers survive two-sided moral hazard, that is, buyers who “borrow” the items for free and abuse the guarantees? This suggests interesting questions for future research. In particular, it points to the evaluate of field experiments that alter guarantee policies and prices to measure more directly the effects on markets, and to use subjects who have experience with guarantee policies.

In sum, this paper illustrates that markets that may be handicapped by moral hazard can introduce simple innovations, such as satisfaction guaranteed, that engage with moral preferences to increase economic efficiency.

References

- Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. *The quarterly journal of economics*, pages 488–500.
- Andreoni, J., Brown, P. M., and Vesterlund, L. (2002). What makes an allocation fair? some experimental evidence. *Games and Economic Behavior*, 40(1):1–24.
- Andreoni, J., Castillo, M., and Petrie, R. (2003a). What do bargainers' preferences look like? experiments with a convex ultimatum game. *American Economic Review*, pages 672–685.
- Andreoni, J., Harbaugh, W., and Vesterlund, L. (2003b). The carrot or the stick: Rewards, punishments, and cooperation. *American Economic Review*, pages 893–902.
- Andreoni, J., Kuhn, M., and Samuelson, L. (2017). Building rational cooperation on their own: Learning to start small. *Working paper*.
- Andreoni, J. and Miller, J. (2002). Giving according to garp: An experimental test of the consistency of preferences for altruism. *Econometrica*, 70(2):737–753.
- Andreoni, J. and Samuelson, L. (2006). Building rational cooperation. *Journal of Economic Theory*, 127(1):117–154.
- Ashraf, N., Bohnet, I., and Piankov, N. (2016). Decomposing trust and trustworthiness. *Experimental Economics*, 9:193–208.
- Barr, A. and Serneels, P. (2009). Reciprocity in the workplace. *Experimental Economics*, 12(1):99–112.
- Berg, J., Dickhaut, J., and McCabe, K. (1995). Trust, reciprocity, and social history. *Games and economic behavior*, 10(1):122–142.
- Bohnet, I., Frey, B. S., and Huck, S. (2001). More order with less law: On contract enforcement, trust, and crowding. In *American Political Science Association*, volume 95, pages 131–144. Cambridge Univ Press.
- Bohnet, I. and Huck, S. (2004). Repetition and reputation: Implications for trust and trustworthiness when institutions change. *American Economic Review*, pages 362–366.
- Bohnet, I. and Zeckhauser, R. (2004). Trust, risk and betrayal. *Journal of Economic Behavior & Organization*, 55(4):467–484.
- Bolton, G. E., Katok, E., and Ockenfels, A. (2004). How effective are electronic reputation mechanisms? an experimental investigation. *Management science*, 50(11):1587–1602.
- Bolton, G. E. and Ockenfels, A. (2000). Erc: A theory of equity, reciprocity, and competition. *American economic review*, pages 166–193.
- Bracht, J. and Feltovich, N. (2008). Efficiency in the trust game: an experimental study of precommitment. *International Journal of Game Theory*, 37(1):39–72.

- Bracht, J. and Feltovich, N. (2009). Whatever you say, your reputation precedes you: Observation and cheap talk in the trust game. *Journal of Public Economics*, 93(9):1036–1044.
- Brown, M., Falk, A., and Fehr, E. (2004). Relational contracts and the nature of market interactions. *Econometrica*, 72(3):747–780.
- Charness, G. and Dufwenberg, M. (2006). Promises and partnership. *Econometrica*, 74(6):1579–1601.
- Charness, G. and Dufwenberg, M. (2010). Bare promises: An experiment. *Economics Letters*, 107(2):281–283.
- Charness, G. and Haruvy, E. (2002). Altruism, equity, and reciprocity in a gift-exchange experiment: an encompassing approach. *Games and Economic Behavior*, 40(2):203–231.
- Charness, G. and Rabin, M. (2002). Understanding social preferences with simple tests. *Quarterly journal of Economics*, pages 817–869.
- Che, Y.-K. (1996). Customer return policies for experience goods. *The Journal of Industrial Economics*, pages 17–24.
- Cooper, R. and Ross, T. W. (1985). Product warranties and double moral hazard. *The RAND Journal of Economics*, pages 103–113.
- Cox, J. C. (2004). How to identify trust and reciprocity. *Games and Economic Behavior*, 46(2):260–281.
- Dufwenberg, M. and Kirchsteiger, G. (2004). A theory of sequential reciprocity. *Games and economic behavior*, 47(2):268–298.
- Durlauf, S. N. (2002). On the empirics of social capital. *The economic journal*, 112(483):F459–F479.
- Eckel, C. C. and Wilson, R. K. (2004). Is trust a risky decision? *Journal of Economic Behavior & Organization*, 55(4):447–465.
- Ellingsen, T. and Johannesson, M. (2004). Promises, threats and fairness. *The Economic Journal*, 114(495):397–420.
- Ellingsen, T., Johannesson, M., Tjøtta, S., and Torsvik, G. (2010). Testing guilt aversion. *Games and Economic Behavior*, 68(1):95–107.
- Engle-Warnick, J. and Slonim, R. L. (2004). The evolution of strategies in a repeated trust game. *Journal of Economic Behavior & Organization*, 55(4):553–573.
- Ensminger, J. (2004). Market integration and fairness: evidence from ultimatum, dictator, and public goods experiments in east africa. *Foundations of human sociality: economic experiments and ethnographic evidence from 15 small-scale societies*, pages 356–381.
- Falk, A., Fehr, E., and Fischbacher, U. (2008). Testing theories of fairness intentions matter. *Games and Economic Behavior*, 62(1):287–303.
- Falk, A. and Fischbacher, U. (2006). A theory of reciprocity. *Games and economic behavior*,

- 54(2):293–315.
- Fehr, E., Gächter, S., and Kirchsteiger, G. (1997). Reciprocity as a contract enforcement device: Experimental evidence. *Econometrica: journal of the Econometric Society*, pages 833–860.
- Fehr, E., Kirchler, E., Weichbold, A., and Gächter, S. (1998). When social norms overpower competition: Gift exchange in experimental labor markets. *Journal of Labor economics*, 16(2):324–351.
- Fehr, E., Kirchsteiger, G., and Riedl, A. (1993). Does fairness prevent market clearing? an experimental investigation. *The Quarterly Journal of Economics*, pages 437–459.
- Fehr, E., Klein, A., and Schmidt, K. M. (2007). Fairness and contract design. *Econometrica*, 75(1):121–154.
- Fehr, E. and List, J. A. (2004). The hidden costs and returns of incentive trust and trustworthiness among ceos. *Journal of the European Economic Association*, 2(5):743–771.
- Fehr, E. and Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *Quarterly journal of Economics*, pages 817–868.
- Glaeser, E. L., Laibson, D. I., Scheinkman, J. A., and Soutter, C. L. (2000). Measuring trust. *Quarterly Journal of Economics*, pages 811–846.
- Gneezy, U. (2005). Deception: The role of consequences. *American Economic Review*, pages 384–394.
- Gneezy, U., Güth, W., and Verboven, F. (2000). Presents or investments? an experimental analysis. *Journal of Economic Psychology*, 21(5):481–493.
- Gneezy, U. and List, J. A. (2006). Putting behavioral economics to work: Testing for gift exchange in labor markets using field experiments. *Econometrica*, 74(5):1365–1384.
- Gneezy, U., Rockenbach, B., and Serra-Garcia, M. (2013). Measuring lying aversion. *Journal of Economic Behavior & Organization*, 93:293–300.
- Gneezy, U. and Rustichini, A. (2000). A fine is a price. *The Journal of Legal Studies*, 29:1–17.
- Güth, W., Schmittberger, R., and Schwarze, B. (1982). An experimental analysis of ultimatum bargaining. *Journal of economic behavior & organization*, 3(4):367–388.
- Heal, G. (1977). Guarantees and risk-sharing. *The Review of Economic Studies*, pages 549–560.
- Heiman, A., McWilliams, B., Zhao, J., and Zilberman, D. (2002). Valuation and management of money-back guarantee options. *Journal of retailing*, 78(3):193–205.
- Hurkens, S. and Kartik, N. (2009). Would i lie to you? on social preferences and lying aversion. *Experimental Economics*, 12(2):180–192.
- Jacobson, S. and Petrie, R. (2013). Favor trading in public good provision. *Experimental Economics*, pages 1–22.

- Janakiraman, N., Syrdal, H. A., and Freling, R. (2016). The effect of return policy leniency on consumer purchase and return decisions: A meta-analytic review. *Journal of Retailing*, 92(2):226–235.
- Keser, C. (2002). Trust and reputation building in e-commerce. *unpublished manuscript*.
- Kessler, A. S. and Lülfesmann, C. (2004). Bilateral bargaining, unverifiable quality, and options to return. *Economic Theory*, 23(2):395–410.
- Knack, S. and Keefer, P. (1997). Does social capital have an economic payoff? a cross-country investigation. *The Quarterly journal of economics*, pages 1251–1288.
- Kube, S., Maréchal, M. A., and Puppe, C. (2013). Do wage cuts damage work morale? evidence from a natural field experiment. *Journal of the European Economic Association*, 11(4):853–870.
- Levine, D. K. (1998). Modeling altruism and spitefulness in experiments. *Review of economic dynamics*, 1(3):593–622.
- List, J. A. (2006). The behavioralist meets the market: Measuring social preferences and reputation effects in actual transactions. *Journal of Political Economy*, 114(1).
- Mann, D. P. and Wissink, J. P. (1988). Money-back contracts with double moral hazard. *The RAND Journal of Economics*, pages 285–292.
- Mann, D. P. and Wissink, J. P. (1990). Money-back warranties vs. replacement warranties: A simple comparison. *The American Economic Review*, pages 432–436.
- Mayzlin, D., Dover, Y., and Chevalier, J. A. (2014). Promotional reviews: An empirical investigation of online review manipulation. *The American Economic Review*, 104(8):2421–2445.
- Pelligra, V. (2011). Empathy, guilt-aversion, and patterns of reciprocity. *Journal of Neuroscience, Psychology, and Economics*, 4(3):161.
- Rabin, M. (1993). Incorporating fairness into game theory and economics. *The American economic review*, pages 1281–1302.
- Rubinstein, A. (1982). Perfect equilibrium in a bargaining model. *Econometrica: Journal of the Econometric Society*, pages 97–109.
- Sobel, J. (2005). Interdependent preferences and reciprocity. *Journal of Economic Literature*, pages 392–436.
- Van Huyck, J. B., Battalio, R. C., and Walters, M. F. (1995). Commitment versus discretion in the peasant-dictator game. *Games and Economic Behavior*, 10(1):143–170.
- Wooldridge, J. M. (2010). *Econometric analysis of cross section and panel data*. MIT press.
- Zak, P. J. and Knack, S. (2001). Trust and growth. *The economic journal*, 111(470):295–321.