Econ. Theory **16**, 719–734 (2000)

Economic Theory

Theoretically robust but empirically invalid? An experimental investigation into tax equivalence*

Rudolf Kerschbamer and Georg Kirchsteiger

Department of Economics, University of Vienna, Hohenstaufeng.9, 1010 Vienna, AUSTRIA
(e-mails:rudolf.kerschbamer@univie.ac.at; georg.kirchsteiger@univie.ac.at)

Reveived: April 30, 1999; revised version: November 15, 1999

Summary. The idea that the final distribution of the tax burden (economic incidence) does not depend on the initial distribution of tax liabilities (statutory incidence) is referred to as the Liability Side Equivalence principle. This paper tests this principle in the laboratory and finds that subjects who actually have to pay the tax carry a higher tax burden. It is argued that this violation of Liability Side Equivalence is due to the fact that a change in the distribution of tax liabilities induces a shift in behaviorally relevant social norms. This shift, in turn, affects the impact of the tax. Our results explain some striking empirical observations and have important theoretical and practical implications.

Keywords and Phrases: Tax equivalence, Tax incidence, Social norms.

JEL Classification Numbers: H21, H22, H30, C91, C92.

1 Introduction

"...it is a matter of indifference whether a general tax on transactions is assessed on the seller's or on the buyer's side of the market" (Musgrave 1959, p. 351). This principle, which we refer to as Liability Side Equivalence (LSE), underlies virtually any theoretical and empirical investigation on the effects of taxes. LSE is theoretically very robust. It holds independently of the nature of the transaction upon which the tax is imposed and independently of the framework in which the

* We are indebted to Gerhard Clemenz, Uwe Dulleck, Simon Gächter, Stefan Krasa and Clemens Puppe for helpful comments. Special thanks for useful suggestions are due to an anonymous referee and to the editor of this special issue. We also gratefully acknowledge financial support from the Austrian Chamber of Labor. Of course, the usual disclaimer applies.
Correspondence to: R. Kerschbamer

transaction is conducted (unilateral or bilateral monopoly, oligopoly, or perfect competition), provided that prices can adjust and that the behavior of the parties in the transaction depends on net payoffs only.[1,2]

Despite its theoretical robustness, policy makers and other economic agents do not act in accordance with the normative implications of LSE. For example, if LSE holds, the policy decision about the legal distribution of tax liabilities should be guided exclusively by the objective to minimize administration and compliance costs. Minimization of these costs would, in general, require imposing the tax liability on one side of the transaction only. Nevertheless, most tax codes distinguish between an employer's and an employee's contribution to the social security/payroll tax. Similarly, reductions of labor taxes imposed on producers' side are frequently discussed as remedies against unemployment. At the same time taxes on the same factor inserted on the household side are rarely mentioned.[3] Apparently, public policy discussions proceed on the assumption that the former are mainly borne by producers while the latter impose few if any burden on them – a belief that is obviously inconsistent with LSE.

The discrepancy between economic theory on the one side and actual tax policy and public discussion of policy issues on the other raises the question of LSE's empirical validity. An empirical test of LSE with field data seems difficult since in the real world economic variables typically depend on a variety of continuously changing factors and not only on taxes.[4] An alternative way of testing LSE are experiments in the laboratory. This method has the advantage that the observations are made under circumstances controlled by the researcher.

In perfectly competitive markets with price taking agents, there is little freedom in price determination. It is therefore not surprising that the policy discussions inconsistent with LSE concentrate on markets characterized by frictions, heterogeneous goods and bargaining. One possible reason for a violation of LSE in such markets is price rigidity. For example, if after a change in tax code the entire social security tax has to be paid by the employers, it is unlikely that gross wages would fall immediately. The labor market would be out of equilibrium in the short run and firms would have to bear a larger share of the tax burden as before. However, even in labor markets, markets with pronounced institutional impediments to price flexibility, wages are adjusted quite regularly. Hence, wage

[1] Cf., for example, Musgrave (1959), Kotlikoff and Summers (1987), and Pisauro (1991).

[2] Several other equivalencies between taxes (e.g. the equivalence between a uniform tax on wages and a uniform tax on consumption, or the equivalence between a uniform tax on income and a uniform tax on output) have been recognized in the public finance literature. These equivalencies hold, however, only under fairly restrictive conditions and are not the focus of our paper.

[3] See e.g. the " Programm für mehr Wachstum und Beschäftigung" (program for more growth and employment) of the German Government (1996).

[4] Cf., however, Lockwood and Manning (1993). These authors analyze the effects of a non-linear tax system for wage bargaining in unionized economies and then test their theoretical predictions empirically. In the empirical investigation they find that income taxes (incident on households) and payroll taxes (incident on firms) have very different long run effects. They argue that "...this presents something of a problem, as any theory based on maximization behavior...would suggest that the formal incidence of a tax should be irrelevant in the long run" (p. 20), and attribute this result to either data problems or model misspecification.

rigidity alone can cause only short run violations of LSE. In addition, changes in the tax system can be foreseen very often by both, trade unions and employers. Thus, wages could be adjusted accordingly even in advance.

An alternative explanation for violations of LSE in not perfectly competitive markets starts from the assumption that the behavior of economic agents is not entirely based on their net gains. For example, it seems quite plausible that the legal obligation to pay a tax is regarded as a moral obligation to bear it (to a certain degree) also economically. Whenever this is the case, individual behavior is affected by gross gains and a shift in the legal liability to pay a tax may influence its economic outcome.

The present paper tests this second explanation in the laboratory. As mentioned earlier, violations of LSE seem to be observed in markets that are not perfectly competitive but rather characterized by bargaining. A very simple bargaining situation, which is also easy to implement in an experiment, is the one stage ultimatum game. In this game one person, the proposer, makes an offer about the division of a given amount of money (the “cake”) between her and her partner, the responder. The responder may accept or reject this offer. If he accepts, the proposer receives the cake minus the offer as gross earnings. The responder’s gross earnings are the offer. If the offer is rejected both players receive nothing.

Assuming rational subjects whose decisions are based on their own net earnings only, the subgame perfect equilibrium of this game is straightforward. Since in case of rejection the responder earns nothing, he is prepared to accept any offer that gives him a nonnegative gain. Knowing this, the proposer will offer zero, and this offer will be accepted by the responder. Thus, in subgame perfect equilibrium the proposer receives the whole cake and the responder earns nothing.[5] It is well known that in one stage ultimatum experiments subjects do not behave in accordance with the subgame perfect equilibrium [see Güth et al. (1982), and Roth et al. (1991), among many others; for an overview see Camerer and Thaler (1995), or Güth (1995)]. The net offers differ substantially from zero, and low but positive offers are frequently rejected. Most authors explain these results by the impact of social norms (or manners) on individual behavior. Subjects have a propensity to punish those who offer them an amount that is, according to their fairness norm, too low, even if punishment is costly for them [see e.g. Bolton (1991), Kirchsteiger (1994), or Camerer and Thaler (1995)].[6] This implies that there exists a threshold such that the offer is rejected if it is below this threshold.

[5] When the proposer can choose his offer continuously then this is the unique sugame perfect Nash equilibrium. Whit a discrete stategy space there exists a second subgame perfect equilibrium in which the responder accepts any strictly positive offer and rejects zero, and in which the proposer offers the smallest monetary unit allowed.

[6] Other explanations are adaptive learning (see Roth and Erev, 1995; Gale et al., 1995) and the “lack of anonymity” hypothesis, i.e. the assertion that the very act of observing behavior by the experimenter influences play away from the subgame perfect equilibrium toward a “fair” allocation (see Hoffman et al., 1994). Abbink et al. (1996) designed an experiment to distinguish between adaptive learning and punishment hypothesis. They found much more evidence for the latter than for the former. Bolton and Zwick (1995) conducted experiments that discriminate between the anonymity and the punishment explanation. They present evidence supporting the conclusion that the punishment

In order to test for LSE, we modify the simple ultimatum game such that in case of acceptance one of the partners has to pay a tax from her/his gross earnings. In half of the sessions the proposer has to pay it, in the other half the responder. In each session subjects play the simple ultimatum game several times (several "rounds") with different partners. We conduct our experiment in two versions. In the "Non Insurance" ("N") version the subjects are divided equally into the groups P (for proposer) and R (for responder), and each subject gets a partner each round. In this version we give subjects no justification for the collection of taxes. In the "Insurance" ("I") version the R group consists of two subjects more than the P group. Hence, two responders have no partner in each round. These subjects are not allowed to participate in the respective round and they get a monetary compensation for that. Now the collection of taxes is motivated with the argument that funds are needed to finance the compensations to inactive players – in the Insurance version the responders were "insured" against not being matched.

Assuming rational individuals whose decisions depend only on their own net earnings the subgame perfect equilibrium prediction of all designs of our experiment is that the responder earns nothing and the proposer receives the cake minus the tax, regardless of upon whom the tax is levied.[7] Hence, LSE should hold if subjects are fully rational and egoistic. On the other hand, if there exists a threshold such that offers below this threshold are rejected, the validity of LSE depends crucially on the nature of the social norm underlying the threshold. As long as this threshold (measured in net earnings) remains unaffected by a change in the statutory distribution of tax liabilities, LSE should hold. If, however, the legal obligation to pay a tax is regarded as a moral obligation to bear it, the threshold should change with a change in the statutory distribution of tax liabilities and LSE should be violated.

We found that LSE is violated in both versions of our experiment. Subjects have, on average, strictly higher after-tax earnings if their partners have to pay the tax than if they have to pay it themselves. Thus, the final distribution of the tax burden depends on the initial distribution of tax liabilities. The reason for conducting the Insurance version was the following. If LSE is violated and if this is due to gross gain based social norms, we thought that the violation of LSE should be more pronounced when people think that there is a plausible reason for taxation. It turned out that the increase in the responders' offered after-tax earnings when the tax liability was shifted from them to the proposers was on average higher in the I than in the N version of the experiment. A statistical test revealed, however, that the difference between the two versions is not significant.

hypothesis explains much more of the deviation from subgame perfect equilibrium than does the anonymity hypothesis.

[7] Since the subjects in the experiment have a discrete strategy space there is actually a second subgame perfect Nash equilibrium outcome in which the responder earns the smallest monetary unit allowed and the proposer the rest (see Footnote 5 above).

The rest of the paper is organized as follows. The next section describes the experimental design. In Section 3 we present the results, in Section 4 we interpret them. Section 5 contains some concluding remarks.

2 Experimental procedures

On the whole we conducted 20 experimental sessions. For each session we recruited – depending on the experimental design (for details see below) – either 10 or 12 prospective subjects. Participants obtained a show up fee of 100 Austrian Shillings (ATS).[8] In addition they received all earnings resulting from their decisions during the experiment. They were paid in cash at the end of the respective session. To rule out interdependencies between different sessions subjects were allowed to participate in a single session only.

At the beginning of each session subjects were randomly allocated to one of two groups. One group was designated to take the role of the proposer, the other group took the role of the responder. Each group was placed in a separate room with one of the experimenters being the only other person present. A total of four to six rounds of the following simple game was played: At the outset of each round the experimenters formed pairs consisting of one proposer and one responder each. To rule out the possibility of reputation formation and of rewarding or punishing a subject's previous behavior, each player obtained a new partner in each round and the identity of the partners was kept secret. Each proposer was asked to make a written proposal on how to divide an amount of 70 ATS between her and her partner, i. e., the responder in the respective pair.[9] The amount of 70 ATS remained the same in all sessions, in all rounds and for all pairs and was common knowledge among subjects. The division proposals ("offers") of all proposers were collected by the experimenter and transmitted to the other experimenter via phone. To rule out any kind of group pressure or herd behavior the transmission was conducted in codified form so that other proposers did not know a given proposer's offer. This codification changed in each round. After having decodified the transmitted offers the experimenter in the responders' room forwarded them to the respective responder. The identity of the person who made the offer was not revealed. Then each responder was asked to decide whether to accept this offer or not. The experimenter in the responders' room collected the acceptance/rejection decisions and transmitted them in codified form via phone to the other experimenter. The latter decodified and forwarded this information in written form to the respective proposer. This concluded the respective round and a new round began or the session was over.

The basic game just described was played in four different designs (see Table 1). Each design was implemented in (all rounds of) 5 of the 20 sessions.

[8] When the experiments were conducted 10 ATS were about 1 US$.

[9] In conducting our experiment we tried to avoid the use of suggestive language that could have an influence on subjects' behavior. For example, we called the partners "Person X" and "Person Y" (and not, for instance, "employer" and "employee"), the offers "division proposals" (and not, e. g., "wage offers"), the payoffs "earnings" (and not, e. g., "profits" and "incomes").....

Table 1. Different designs

	Non insurance version	Insurance version
Tax paid by P	NP	IP
Tax paid by R	NR	IR

Non insurance version: The subjects are divided equally into the groups P and R and each subjet gets a partner each round. Insurance version: The group R consists of two more subjects than the group P: two subjects R remain unmatched; unmatched Rs get a lump sum payment ("insurance" against not being matched).

With respect to the number of participating responders the four designs can be divided in two design-pairs ("versions"). For the **Non Insurance (N)** version of the experiment we recruited 12 students per session. The subjects who showed up on time were divided evenly into the groups P and R, i. e., in this version there were as many proposers as responders.[10] The earnings of the players were determined as follows: If the offer was accepted the gross earnings of the proposer were given by the 70 ATS minus the offer. The gross earnings of the responder were just the offer. If the offer was rejected, both partners received nothing and the 70 ATS expired. In case of acceptance (and only in this case) one of the partners had to pay a tax of 20 ATS from his gross earnings. In 5 of the 10 N sessions the proposer had to pay the tax (we denote this design in what follows by **NP**), in the other 5 N sessions the tax was imposed on the responder (design **NR**).[11] The players got neither a justification for the collection of the tax nor for the choice of the party on which the tax was imposed.

For the **"Insurance" (I)** version of the experiment we recruited 10 students per session. The subjects who showed up on time were again divided into the groups P and R. In contrast to the N version of the experiment the participants were not divided evenly but in such a way that group R had two members more than group P. Hence, in each round two responders were ignored in the process of forming pairs. These two players remained without partner and could not participate in the respective round of the session. The pair formation process was designed in such a way that (i) each active player received a new partner in each round and (ii) each responder remained unmatched for exactly two rounds. Unmatched responders obtained a compensation of 20 ATS in the respective round. The gross earnings of the active players were determined exactly as in the N version of the experiment. In the case of acceptance (and only in this

[10] Unfotunately some subjects who had signed up for the expiriment did not show up on time. If an odd number of subjects came to a session one person was chosen randomly. This person got the show-up fee of 100 ATS but was not allowed to participate.

[11] To rule out uncontrollable side effects caused by queer subjects the range of possible offers in each design was constrained in such a way that each of the two partners in a pair could potentially get net earnings between 0 and 50 ATS. That is, each proposer was asked to choose an offer between 0 and 50 ATS in design NP and an offer between 20 and 70 ATS in design NR. In the experiment no proposer ever tried to make a forbidden offer.

case) one of the partners in each pair had to pay a tax of 20 ATS from his gross earnings. In (all rounds of) 5 of the 10 I sessions the proposer had to pay the tax (we denote this design by **IP**), in the other 5 I sessions the tax was imposed on the responder (design **IR**). The collection of taxes was now motivated with the argument that funds to finance the compensations to the inactive players were needed. Notice, however, that the payments to unmatched responders were pure lump-sum, i.e. did not depend on any decision made by any subject.[12] As in the N version of the experiment the subjects got no motivation for the choice of the "side of the market" on which the tax was imposed.

To make sure that all subjects had effectively understood the rules of the respective session they were asked to compute their own gains and the gains of their partners in four hypothetical examples. After all subjects had solved these examples correctly, and after the experimenters had answered all remaining questions the experiment started with a trial round (i.e., a round without monetary consequences for the players). Then the first round of the "real" game in the respective session began. The number of rounds was such that each person played once with each member of the other group.

3 Experimental results

As described above each design was implemented in 5 of the 20 sessions. Each session lasted about 45 minutes. Each set of students participated in a single session only. Most of the subjects were students of law, some were students of computer sciences and a few were students of psychology. The subjects were comparable in age and education to undergraduates. None of them was a student of the economics department.

In total 48 subjects participated in the NP, 52 subjects in the NR, 46 subjects in the IP and 44 subjects in the IR sessions.[13] Subjects made 116 offers in the NP, 138 offers in the NR, 102 offers in the IP and 93 offers in IR sessions. In the NP sessions 24 offers, in the NR sessions 33 offers, in the IP sessions 22 offers and in the IR sessions 15 offers were rejected. Subjects earned on average 189 ATS (about US$ 19) per session.[14]

---

[12] We decided against making the transfers of unmatched responders dependent on the decisions of matched subjects for the following reason: Since players in ultimatum games seem to be motivated not only by their own monetary (net) payoff but also by their relative payoff standing, this would have introduced additional effects that are absent in the N version of the experiment. (For instance, the desire to help unmatched subjects could have induced a matched responder to accept an offer that otherwise had been rejected.) This would have rendered the interpretation of any differences in the results between the two versions rather difficult. For example, if the violation of LSE turned out to be more pronounced in the I than in the N version of the experiment, it would not be clear whether the reason is the justification of the tax or the additional distributional concerns.

[13] As mentioned before, in some sessions not all subjects enrolled for the experiment actually showed up. For example, in none of the NP sessions more than 10 subjects arrived at the agreed upon date. Hence, in this design a maximum of 5 rounds was played per session.

[14] Note that this amount exceeds by far the hourly net wage an average Austrian student could earn in his best alternative job (about 80 ATS).

The mean offered net-income for the responder (i.e., offer net of taxes) over all sessions was 19.3, i.e. a little less than 40% of the net cake. This replicates roughly the results of previous studies (see Camerer and Thaler, 1995). However, the offers differed in the different designs as our first result shows. In this result reference is made to the term *net offer*. By net offer we mean the after-tax earnings of the responder implied by a given offer. That is, the net offer equals the actual offer in the NP and the IP design, while it is the offer minus the tax in the NR and the IR design.

Result 1. *In both versions of our experiment the net offers were, on average, higher if the tax was imposed on the proposer rather than on the responder.*

Figure 1 presents evidence for this finding. The figure compares the distributions of all net offers in the two designs belonging to the same (i.e., N or I) version of the experiment.

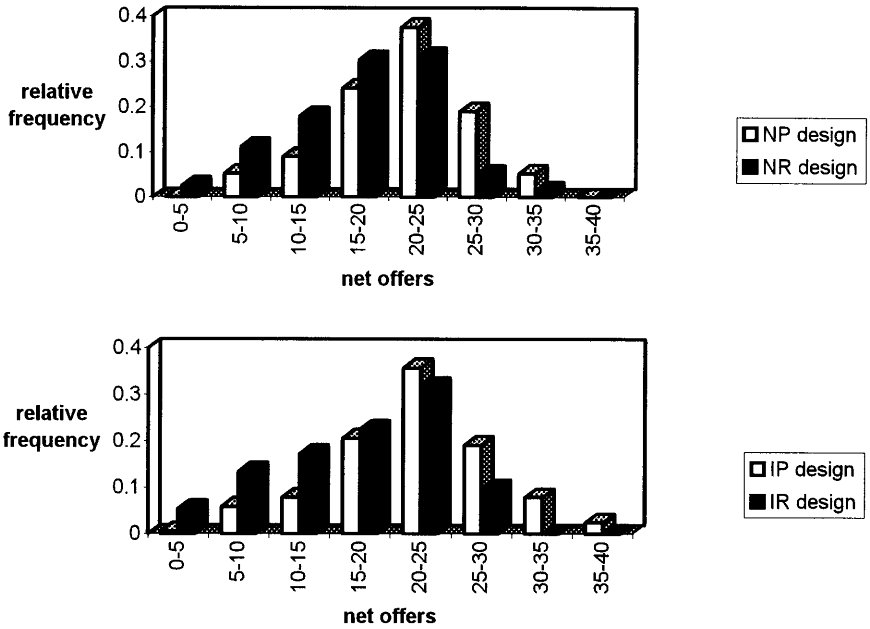

Figure 1. Distribution of net offers

It is obvious from Figure 1 that the relative frequency of low net offers (offers less than 20) was higher in the NR design (IR design) than in the NP design (IP design), whereas high net offers were made more often in the NP (IP) than in the NR design (IR design). Hence it was, on average, always better for a responder if the tax was imposed on the proposer rather than on the responder.

The same conclusion can be drawn from Table 2. In the N version of our experiment the average net offer was 21.1 ATS if the proposer had to pay the tax, whereas the net offer was only 17.6 ATS if the responder had to pay the tax, implying a difference of 3.5 ATS, i.e. 17.5% of the tax. For net earnings this implies that the responder's offered net income increased about 20 percent when the tax liability was shifted from him to the proposer. The violation of LSE was even more pronounced in the I version of the experiment. Here the average difference in net offers was 4.8 ATS, i.e. 24% of the tax. This implies that in the I version the responder's offered net income increased by about 25% if the tax liability was shifted from him to his partner.[15]

Table 2. Average net offers

Design	All net offers	Last round net offers
NP	21.1	21.4
NR	17.6	17.4
NP minus NR	3.5	4.0
IP	21.8	21.9
IR	17.0	13.9
IP minus IR	4.8	8.0

Next we tested the equivalence statistically. Since the same subjects participated in the different rounds of a given session individual observations are not independent from each other. Hence, only individual first round observations and mean offers of the different sessions were used for statistical tests. We first conducted a Wilcoxon-Mann-Whitney test with individual first round offers. For both versions the differences in net offers were statistically significant even at a 1% level.[16,17] Furthermore, we conducted a permutation test with average offers of each session as well as with the average last round offers.[18]

Table 3 shows the significance levels for the hypothesis that the session means of all net offers (or the session means of last round net offers, respectively) in the NP sessions (IP sessions) are equal to the session means of all net offers (session means of last round net offers) in the NR sessions (IR sessions). The alternative hypothesis is that the mean offers are higher in the NP sessions (IP sessions) than in the NR sessions (IR sessions). As can be seen from Table 3 the hypothesis that the session averages in the NP sessions are equal to those in the NR sessions can be rejected at the 2% confidence level. The same holds for comparisons within the I version of our experiment.

[15] We also conducted statistical tests on whether the violation of LSE is larger in the I than in the N version of the experiment. The difference turned out not to be significant, however.

[16] We also conducted a Kolmorogov- Smirnov test with the individual first round offers. The results are qualitatively the same as those of the Wilcoxon-Mann-Whitney test.

[17] We tested the equivalence also by using all individual observations. The results of these Wilcoxon-Mann-Whitney tests are qualitatively the same as those for the individual first round offers. However, as already mentioned, the individual observations were not independent from each other. Hence, tests using individual observations are not statistically valid.

[18] Since we have numerical data, the most powerful nonparametric test of location is the permutation test. However, this test is only feasible when the sample size is small. If we use all individual first round offers the sample size is too large to make a permutation test. Hence, in this case we have to conduct a Wilcoxon-Mann-Whitney test. On the other hand, the use of session means allows for the

Table 3. Results of the permutation tests

	p (NP vs. NR)	p (IP vs. IR)
Session mean of all net offers	0.012	0.012
Session mean of last round net offers	0.031	0.004

p (A vs. B): Significance level of the permutation test statistics with data of the A and the B sessions

The differences between the P and R designs within each version of the experiment did not vanish when the subjects got more experienced. On the contrary, the average differences in the net offers were even more pronounced in the last rounds (see Table 2). These last round differences were highly significant (see Table 3). Figure 2 shows the evolution of the differences in net offers. It is not clear whether these differences converged to any specific level, but obviously there was no tendency for the differences to vanish.

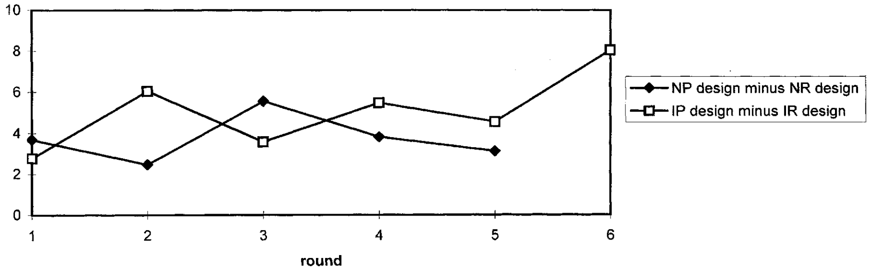

Figure 2. Differences between the mean net offers in different rounds

The failure of LSE is also reflected in the acceptance decisions of the responders, as our second result shows:

Result 2. *When controlling for the net amount offered the rejection rates were higher if the tax was imposed on the proposer rather than on the responder. This result holds for both, the Insurance and the Non Insurance version of the experiment.*

Figure 3 depicts the rejection rates following net offers of different levels. As can be seen from this figure, net offers of a given level were more often rejected if the tax was imposed on the proposer. In other words, if the responder had to pay the tax he was more willing to accept a given after-tax income than if the proposer had to pay it.[19] This holds for all levels of net offers and for both versions of the experiment. A similar conclusion can be also drawn from Table 4. This table shows the rejection rates of net offers below 20.

permutation test. For a discussion of the applicability of the differnt tests see Siegel and Castellan (1988, cap. 6).

[19] In the NP design no offer below 5 was ever made and in the NP as well as in the NR design all offers above 25 were accepted. In the Insurance version no offer above 30 was rejected.

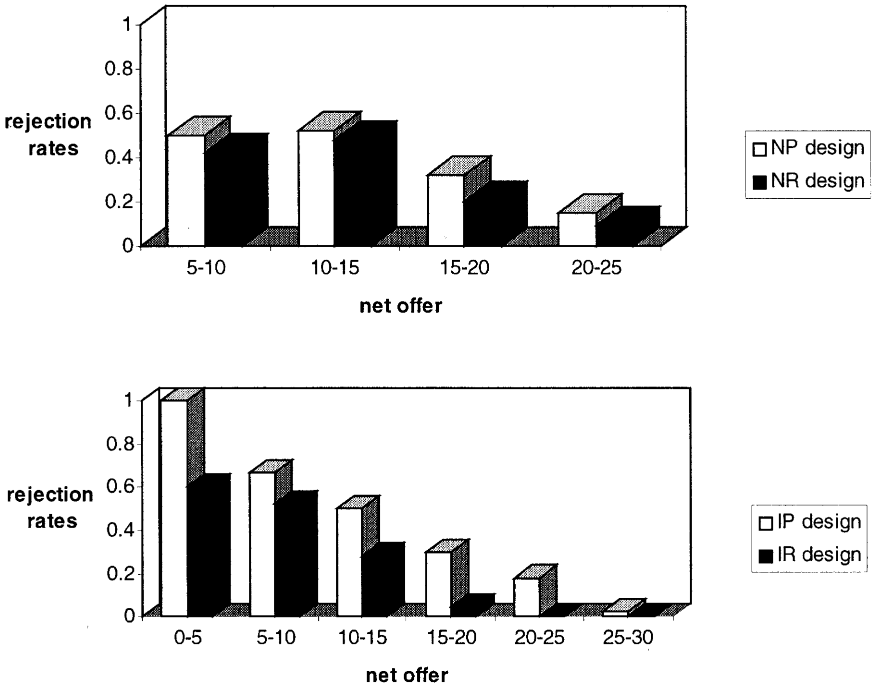

Figure 3. Rejection rates

Again these rates were higher when the tax was imposed on the proposer. To see whether these differences were significant we conducted a permutation test with the rejection rates of offers below 20 in each individual session. It turns out that the differences were significant at a 10% level in the Non Insurance version and at a 5% level in the Insurance version. Hence, we can conclude that the acceptance behavior of the responders – like the offers made by the proposers – contradicted LSE. We can therefore summarize our results to:

Statement: *Liability Side Equivalence held neither in the Non Insurance nor in the Insurance version of our experiment. In both versions subjects had, on average, strictly higher net earnings if the tax was imposed on their partners than if they had to pay the tax themselves.*

Table 4. Rejection rates of net offers below 20

NP	NR	IP	IR
0.55	0.43	0.57	0.36

4 Interpretation of the results

Having established the fact that LSE is violated in the laboratory, one has to ask for the reasons for this result. One possible explanation is, of course, money illusion. If subjects suffer from money illusion they accept lower net offers in settings in which they are asked to pay a tax because they do not really realize that taxes reduce their net earnings. However, we do not believe that money illusion explains our results, since in our experiment (i) there was no time lag between disbursement of gross earnings and payment of taxes, and since (ii) subjects repeatedly revealed their ability to calculate net earnings of both parties correctly. We believe that the violation of LSE can better be explained by the influence of social norms on the behavior of players. Recall the punishment explanation for behavioral anomalies in ultimatum experiments without taxes. In this explanation people are willing to reject strictly positive offers below a certain threshold because they consider such offers as unfair. We think that a change in the distribution of statutory liabilities to pay a tax induces a shift in subjects' fairness perceptions. This shifts subjects' acceptance thresholds and thus leads to violations of LSE.[20,21]

To facilitate the understanding of this explanation for our results it might be instructive to construct a simple model that incorporates social norms and that is consistent with the violation of LSE. Consider a game with two players indexed by $i \in \{P, R\}$ and let x^i denote the monetary net payoff of, and t^i the tax paid by, player i. Assume that each player acts to maximize the expected value of his/her utility function

$$u^i = u^i \left(x^i, s^i \right), \qquad (1)$$

where $s^i \equiv x^j - x^i + \lambda \left(t^j - t^i \right)$ and $\lambda \in [0, 1]$. We assume that

(i) u^i is continuously differentiable on the domain of $\left(x^i, s^i \right)$;
(ii) $u_1^i > 0$;
(iii) $u_2^i \leq 0$ if $s^i \geq 0$, and $u_2^i < u_1^i/2$ if $s^1 < 0$.

The function specified in (1) is a statement about the objectives that motivate behavior of economic agents. It tells us that the behavior of subjects might be driven not only by their own material wellbeing, but also by the extend to which their social norm concerning the distribution of payoffs is fulfilled. To evaluate the implications of (1) for the outcome of our ultimatum game let us start with the special case where $\lambda = 0$. In this case the social norm depends on net payoffs

[20] This explanation is, in most respects, empirically indistinguishable from money illusion. However, for most policy questions it is a matter of indifference which of these two explanations is the more appropriate one.

[21] As mentioned earlier there exist other explanations for the ultimatum game results, an important one being the adaptive learning hypothesis. These alternative explanations are incompatible with our results as long as they are based on the assumption that subjects' behavior depends only on net earnings. We regard it as implausible that learning dynamics depend on gross earnings. If, however, ultimatum experiment results are really caused by learning, our results show that learning behavior is (at least partly) based on gross gains and that LSE does not hold as long as the learning dynamics do not converge to the subgame perfect equilibrium.

only. If the net payoff of player j exceeds that of player i $\left(s^i > 0\right)$ then $u_2^i \leq 0$ measures the utility loss from being worse off: For a given own material net payoff x^i player i prefers an outcome with more equal material payoffs to one in which the inequality in his opponent's favor is more pronounced. For the case in which player i has a higher monetary payoff than player j $\left(s^i < 0\right)$ we allow a wide variety of possible preferences. Player i might be spiteful and prefer a further increase in inequality $\left(u_2^i < 0\right)$ as in Kirchsteiger (1994), he may be driven by fairness/equity considerations and prefer a more equal distribution $\left(u_2^i > 0\right)$ as in Bolton and Ockenfels (1999), Fehr, Kirchsteiger and Riedl (1998), and Fehr and Schmidt (1999), or he may not care about inequality in his own favor at all $\left(u_2^i = 0\right)$. All we require is that fairness considerations do not overcome the direct impact of the own material net payoff $\left(u_2^i < u_1^i/2\right)$.

Obviously, if $\lambda = 0$, LSE holds. Also notice that, if $u_2^i = 0$ for all s^i, we are back to the standard model in which players are only interested in their own material net payoffs.

Now suppose that the legal obligation to pay a tax is regarded as a moral obligation to bear it (to a certain degree) also economically ($\lambda > 0$). Also suppose that subjects exhibit a strict aversion against being worse off $\left(u_2^i < 0 \text{ if } s^i > 0\right)$. Then a player's wellbeing does not only depend on the distribution of net- but also on that of gross- earnings. (In the limit case of $\lambda = 1, s^i$ gives weight only to the distribution of gross-earnings). What does this imply for the outcome of our ultimatum game, where, in case of an agreement, the proposer has to pay a tax of t^P and the responder one of t^R? To answer this question let the proposer's preferences be represented by $u^P = u^P\left(x^P, s^P\right)$ and assume that the proposer does not exactly know the preferences $u^R = u^R\left(x^R, s^R\right)$ of a randomly matched responder, but only the distribution of responders' preferences. Denote by c the net cake to be divided, and assume that $t^i(i = P, R)$ is lower than c, as it is the case in our experiment. First notice that, if the players knew each others preferences, then we would never observe rejected offers. To see this, look at the allocation where, for given taxes, the difference in the weighted average of gross and net earnings is zero $\left(s^i = 0\right)$. For this allocation the utilities of proposer and responder are given by $u^P\left(\left(c + \lambda\left(t^R - t^P\right)\right)/2, 0\right)$ and $u^R\left(\left(c - \lambda\left(t^R - t^P\right)\right)/2, 0\right)$, respectively. Since the impact of the own monetary net payoff on utility is strictly positive, both players strictly prefer this allocation to the disagreement allocation which yields utility $u^i = u^i(0, 0)$ for $i = P, R$.

Next notice that, again if players' preferences were common knowledge, a proposer would always offer the lowest amount accepted by the responder. This follows from $u_2^i < u_1^i/2$. By the same assumption, for every responder there exists a unique threshold such that he accepts any offer higher or equal, and rejects any offer lower, than this critical one. For the case where preferences are private information these two facts together imply that a proposer's behavior does not directly depend on the distribution of responders' preferences but only on the distribution of acceptance thresholds. Now consider a given responder R. Denote his acceptance threshold by $\underline{x}^R$. Since the responder accepts an offer only if the utility from acceptance is not lower than the disagreement utility, $\underline{x}^R$

solves

$$u^R(0,0) = u^R\left(x^R, c - 2x^R + \lambda\left(t^P - t^R\right)\right). \tag{2}$$

For $x^R = 0$ the RHS of (2) is smaller than $u^R(0,0)$. Also, for $x^R = \frac{c+\lambda\left(t^P - t^R\right)}{2}$ the RHS of (2) exceeds $u^R(0,0)$ as has been shown above. Therefore, and because u^R is continuous and strictly monotone in x^R, a unique solution of (2) exists, and this solution $\underline{x}^R$ is smaller than $\left(c + \lambda\left(t^P - t^R\right)\right)/2$. Thus, $s^R > 0$ and $u_2^R < 0$ at $\underline{x}^R$. Furthermore, $\underline{x}^R$ depends on λ, t^P and t^R. Applying the implicit function theorem reveals that

$$\begin{aligned} \partial\underline{x}^R/\partial t^P &= -\lambda u_2^R/\left(u_1^R - 2u_2^R\right) > 0; \\ \partial\underline{x}^R/\partial t^R &= \lambda u_2^R/\left(u_1^R - 2u_2^R\right) < 0. \end{aligned}$$

In other words, a shift of the formal obligation to pay a tax from the proposer to the responder decreases the responder's acceptance threshold. Since this is true for all responders in the population, the distribution of acceptance thresholds for the case where $t^P = \bar{t} > 0$ and $t^R = 0$ dominates that for the case where $t^P = 0$ and $t^R = \bar{t}$ in the sense of first order stochastic dominance. The rest is trivial: Higher acceptance thresholds in the population of responders if the tax is imposed on proposers induce proposers to make more generous net offers (our Result 1), and responders to reject net offers of a given level more frequently (our Result 2). This violates LSE. Furthermore, the violation of LSE increases in λ. This might explain why the departure from LSE seems to be more pronounced in the I than in the N version of our experiment: Subjects' respect for the statuary distribution of the tax liabilities may simply be intensified if they see a reason for collecting taxes.[22]

5 Concluding remarks

The idea that the final distribution of the tax burden (economic incidence) is independent from the initial distribution of tax liabilities (statutory incidence) is referred to as the Liability Side Equivalence (LSE) principle. This paper tested this principle in the laboratory and found that the equivalence was violated. In both versions of our experiment the after-tax earnings offered by the proposers to the responders (net offers) were, on average, higher if the tax was imposed on the proposer than if the responder had to pay it. Furthermore, the responder was more inclined to accept a net offer of a certain level if he had the legal obligation to pay the tax.

[22] This explanation is in line with the finding of Hoffman and Spitzer (1985) and Hoffman et al. (1994) that the results of bargaining depend crucially on the way the bargaining positions are allocated. If there is a justification for a certain allocation of the positions, e.g. if the winners of a quiz "earn" the position of proposers as in Hoffman et al. (1994), unequal divisions of the cake are much more often demanded and accepted. On the importance of entitlements see also Kahneman et al. (1986).

We have demonstrated that a simple model, based on the premise that people are motivated by both their pecuniary after-tax payoffs and social norms concerning the distribution of payoffs, can explain the violation of LSE, provided the legal obligation to pay a tax is regarded as a moral obligation to bear it. LSE relies on two crucial assumptions: (i) People are fully rational. (ii) People are only interested in after-tax earnings.[23] In principle one can give up either of these assumptions in an attempt to explain the evidence. In our model we have chosen to relax assumption (ii) because we belief that the ultimatum game is a relatively simple game and because subjects repeatedly revealed their ability to calculate the after-tax earnings of both parties in the game correctly.

Whether a shift of the legal liability to pay a tax can influence its economic outcome depends, of course, on the environment considered. If trade takes place on competitive markets and if the characteristics of the good traded are completely specified, prices and quantities converge to the market clearing level rather quickly.[24] It seems that in such environments market forces are very strong and overcome any kind of idiosyncrasies of individual behavior.[25] If, on the other hand, prices are determined at least partially by bargaining, the outcome can be shaped by social norms which, in turn, may be influenced by statutory tax incidence. Hence, in such markets there exists the possibility that the outcome is affected by behavior that is influenced by gross gains. An important example for such a market is the labor market. This might explain why there is so much discussion on the distribution of the statutory incidence of the payroll tax but at the same time scarcely a discussion on the statutory incidence of consumption taxes.

The most immediate normative implication of LSE is that the decision over the division of statutory liabilities to pay a tax on the two sides of the market should only be guided by the goal of minimizing administration and compliance costs. This paper's findings suggest that adhering unconditionally to that implication may lead to suboptimal results in markets in which social norms have an influence on equilibrium behavior of transactors: If in such markets a shift in statutory incidence changes transactors' behaviorally relevant norms then the usual trade-off between different basic objectives of tax policy – as e.g., the minimization of the tax induced dead weight loss, the achievement of an equitable distribution of the tax burden, high employment, etc. – must be made in designing the pattern of legal liabilities.

[23] LSE *does not* rest on the assumption that subjects are only interested in their *own* after-tax earnings (selfish behaviour), however.

[24] For an overview of market experiments see Davis and Holt (1993, chapters 3 and 4).

[25] Notice that the kind of motivation or utility function we introduced in Section 4 to explain the violation of LSE in our ultimatum game is consistent with the experimental fact that market games tend to converge to the competitive equilibrium. This has independently been shown by Bolton and Ockenfels (1999) and by Fehr and Schmidt (1999).

References

Abbink, K., Bolton, G., Sadrieh, A., Tang, F.: Adaptive learning versus punishment in ultimatum bargaining, Discussion Paper No. B-381, University of Bonn (1996)

Bolton, G.: A comparative model of bargaining: theory and evidence. American Economic Review **81**, 1096–1136 (1991)

Bolton, G., Ockenfels, A: A theory of equity, reciprocity and competition. American Economic Review (forthcoming) (1999)

Bolton, G., Zwick, R.: Anonymity versus punishment in ultimatum bargaining. Games and Economic Behavior **10**, 95-121 (1995)

Camerer, C., Thaler, R.: Anomalies: ultimatums, dictators and manners. Journal of Economic Perspectives **9**, 209–219 (1995)

Davis, D., Holt, C.: Experimental economics. Princeton, NJ: Princeton University Press 1993

Fehr, E., Schmidt, K.: A theory of fairness, competition and cooperation. Quarterly Journal of Economics (forthcoming) (1999)

Fehr, E., Kirchsteiger, G., Riedl, A.: Gift exchange and reciprocity in competitive experimental markets. European Economic Review **42**, 1–34 (1998)

Gale, J., Binmore, K., Samuelson, L.: Learning to be imperfect: the ultimatum game. Games and Economic Behavior **8**, 56–90 (1995)

German Government: Programm für mehr Wachstum und Beschäftigung. http://www.bundesregierung.de/inland/bpa/themen/t hx059602.htm. (1996)

Güth, W.: On ultimatum bargaining experiments – a personal review. Journal of Economic Behavior and Organisation **27**, 329-344 (1995)

Güth, W., Schmittberger, R., Schwarze, B.: An experimental analysis of ultimatum bargaining. Journal of Economic Behavior and Organisation **3**, 367–388 (1982)

Hoffman, E., Spitzer, M.: Entitlements, rights and fairness: an experimental examination of subjects' concepts of distributive justice. Journal of Legal Studies **15**, 254–297 (1985)

Hoffman, E., McCabe, K., Shachat, K., Smith, V.: Preferences, property rights and anonymity in bargaining games. Games and Economic Behavior **7**, 346–380 (1994)

Kahneman, D., Knetsch, J., Thaler, R.: Fairness as a constraint on profit seeking: entitlements in the market. American Economic Review **76**, 728–741 (1986)

Kirchsteiger, G.: The role of envy in ultimatum games. Journal of Economic Behavior and Organization **25**, 373–390 (1994)

Kotlikoff, L., Summers, L.: Tax incidence. In: Auerbach, A., Feldstein, M., (eds.) Handbook of public economics, Vol.**1**, pp. 1043–1092. Amsterdam: North Holland 1987

Lockwood, B., Manning, A.: Wage setting and the tax system: theory and evidence for the United Kingdom. Journal of Public Economics **52**, 1–29 (1993)

Musgrave, R.: The theory of public finance. New York: McGraw-Hill 1959

Pisauro, G.: The effect of taxes on labor in efficiency wage models. Journal of Public Economics **46**, 329–345 (1991)

Roth, A., Prasnikar, V., Okuno-Fujiwara, M., Zamir, S.: Bargaining and market behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: an experimental study. American Economic Review **81**, 1068–1095 (1991)

Roth, A., Erev, I.: Learning in extensive form games: experimental data and simple dynamic models in the intermediate term. Games and Economic Behavior **8**, 164–212 (1995)

Siegel, S., Castellan, J.: Nonparametric statistics for the behavioral sciences. Singapore: McGraw-Hill 1988