

WORKING PAPERS

Are We Taxing Ourselves? **How Deliberation and Experience Shape Voting on Taxes**

Rupert Sausgruber

Jean-Robert Tyran

October 2010

Working Paper No: 1010



DEPARTMENT OF ECONOMICS

UNIVERSITY OF VIENNA

All our working papers are available at: <http://mailbox.univie.ac.at/papers.econ>

Are We Taxing Ourselves?

How Deliberation and Experience Shape Voting on Taxes

Rupert Sausgruber and Jean-Robert Tyran^{*}

October 2010

We let consumers vote on tax regimes in experimental markets. We test if taxes on sellers are more popular than taxes on consumers, i.e. on voters themselves, even if taxes on sellers are inefficiently high. Taxes on sellers are more popular if voters underestimate the extent of tax shifting in the market. We show that inexperienced voters are prone to such a tax-shifting bias, that experience is an effective de-biasing mechanism, but that pre-vote deliberation about tax regimes makes initially held opinions more extreme rather than correct. Our results suggest that voting on taxes is prone to bias and that easy-to-interpret facts are needed to de-bias voters.

Keywords: Tax shifting, tax liability side equivalence, learning, deliberation, voting

JEL-codes: C92, H22, D72

^{*} Sausgruber: University of Innsbruck, Department of Public Economics, Universitätsstr. 15, A-6020 Innsbruck. Rupert.Sausgruber@uibk.ac.at

Tyran: University of Vienna, Department of Economics, Hohenstaufengasse 9, A-1010 Vienna, Austria, and University of Copenhagen, Department of Economics, Øster Farimagsgade 5, DK-1353 Copenhagen, Denmark. Jean-Robert.Tyran@univie.ac.at

An earlier version of this paper was circulated under the title “Tax salience, voting, and deliberation”. We gratefully acknowledge financial support by the Austrian Science Fund (FWF), under Project No. P17029 and S10307-G14. We thank Wolfgang Höchtel, Alexander Sebald and Frederik Øvlisen for helpful comments.

1 Introduction

A fundamental proposition of public economics claims that it is matter of indifference whether a tax is levied on buyers or sellers. This proposition is called Tax Liability Side Equivalence (Tax LSE for short)¹ and any textbook explains that economic incidence, i.e. who ultimately bears the tax, is independent of legal incidence, i.e. which side of the market formally pays the tax, because taxes are shifted in markets. While Tax LSE is straightforward in theory and is supported by experimental evidence for a broad range of market institutions², it is not obvious to laypeople. Voters may underestimate tax shifting, inducing the erroneous belief that taxing “the other side of the market” comes at a lower economic cost than taxing themselves. Such a “tax-shifting bias” may thus induce voters to believe that economically equivalent taxes are different or, in case of extreme bias, to think that high taxes levied on others are better than low taxes on themselves. Thus, while competitive markets seem to reliably produce outcomes in line with Tax LSE in reaction to exogenously imposed taxation, endogenous choice of such taxes in a vote may not. The tax-shifting bias may well be irrelevant for outcomes in competitive markets but distort democratic choice on taxation.

This paper uses an experimental approach to test for tax-shifting bias and its consequences for voting on tax regimes. We let consumers vote on tax regimes in experimental markets. We test if taxes on sellers are more popular than taxes on consumers, i.e. on voters themselves, even if taxes on sellers are inefficiently high. We show that about a third of inexperienced voters are prone to a tax-shifting bias and that the bias is costly to voters. For example, we find that the median voter is willing to tax the other side of the market at a 25 percent (or 6.4 points) higher rate to avoid a tax levied on themselves, resulting in an income loss to voters of about 20 percent.

A particular focus of our paper is on factors that may de-bias voters, i.e. induce optimal voting. We consider learning from experience and pre-vote deliberation. We find that pre-vote deliberation about tax regimes does not improve optimal voting overall but induces group polarization, i.e. makes initially held opinions more extreme rather than correct. In contrast, we find that experience is an effective de-biasing mechanism. By experience we mean allowing voters to learn either from own or others’ experience. For example, we find that providing voters with clear-cut information about the adverse effects of voting for

1 Tax LSE is also called “invariance of incidence proposition” or “Dalton's Law” in the literature.

2 Tax LSE can be cleanly tested in the laboratory and has been shown to hold for double auction markets (Kachelmeier, et al., 1994), posted offer markets (Borck et al., 2002), and pit markets (Ruffle, 2004). For a review and discussion of this literature see Riedl and Tyran (2005) who find that tax LSE also holds in a gift-exchange labor market.

inefficiently high taxes in other markets reduces the willingness to impose such taxes by about 80 percent (a drop of 5.3 tax points). We also provide suggestive evidence that tax attitudes “in the wild”, i.e. on taxing goods such as electricity or beer, are correlated with voting in the lab, and that providing information about the adverse effects of taxation in the lab induces transfer learning, i.e. brings attitudes “in the wild” more in line with Tax LSE. Our results suggest that voting on taxes is prone to bias and that easy-to-interpret facts are needed to de-bias voters.

Anecdotal evidence from the field suggests that popular perception differs markedly from the textbook. For example, Graetz (2002: 270) claims that “many families underestimate their payroll tax burden because the employers’ share is hidden to employees”, and public debate in the media and the political arena is often much concerned with statutory incidence (see Blinder, 1988 for a discussion and Ruffle, 2004 or Blumkin and Menirav, 2009 for examples). However, it is important to note that a concern for statutory incidence in policy debates is no proof that voters are biased. In fact, statutory incidence does make a difference to economic incidence if prices are rigid and cannot freely adjust to changes in taxation (see Riedl and Tyran, 2005 for a discussion). The relevance of price and wage rigidity is much debated in the field but easy to control in the laboratory. An advantage of our experimental approach is therefore that we can implement market institutions which are known to have flexible prices and quick equilibration, implying that tax LSE holds. Another advantage of an experimental approach is control over information conditions and the ability to isolate causal factors (see Falk and Heckman, 2009 for a discussion). For example, a tax bias is hard to pin down in the field because inattention to in principle known information is often difficult to distinguish from sheer lack of information (see Chetty et al., 2009 for a notable exception).

We proceed in four steps. In a first step, we show that voting on commodity taxes is prone to tax-shifting bias in a laboratory experiment. Participants earn incomes by trading a good in a market. After some periods of trading, consumers are told that market transactions need to be taxed in the future and that tax revenues are wasted, i.e. not returned to market participants. Only consumers can vote, and they vote on implementing a buyer tax or a seller tax. To make the tax-shifting bias costly, consumers are presented with a choice between a buyer tax of t and a seller tax of more than t . According to Tax LSE, consumers correctly perceive that higher taxes reduce their market income independently of which side of the market is taxed and, therefore, vote for imposing a tax on themselves. In contrast, tax-shifting bias induces consumers to vote for an inefficiently high seller tax. In line with the tax-shifting bias hypothesis, we find that consumers vote for the seller tax in 2 out of 3 cases.

In a second step, we investigate if pre-vote deliberation mitigates the tax-shifting bias. We compare voting when voters can freely discuss the two tax regimes with a treatment without this option. We show that deliberation is persuasive, i.e. that the quality of messages a voter receives affects his or her voting. Correct messages tend to improve voting while incorrect messages tend to promote voting for the seller tax. We find that deliberation induces agreement, i.e. that opinions become more homogenous in the presence of deliberation. However, since opinions are divided at the beginning of the deliberation process, voters by chance receive more correct or incorrect messages. As a result, initially held opinions tend to become more extreme and voters agree on more or less arbitrary values (see Glaeser and Sunstein, 2007 for a discussion). We therefore find that deliberation does not de-bias voting of inexperienced voters.

In a third step, we repeat the trading and voting sequence to study if experience mitigates the tax-shifting bias. We find that voters who experienced the income-reducing effects of seller taxes tend to vote for the buyer tax in the sequel. Absent such experience, repeatedly discussing the tax proposal does not reduce tax-shifting bias.

In a fourth step, we run two treatments to address issues of external and internal validity. As a check on external validity, we compare attitudes on taxation of various goods such as electricity or beer with voting behavior in the experiment. We find that tax attitudes in the wild are related to voting in the lab, suggesting that what we measure in the lab has some resemblance to more complex naturally occurring tax issues. As a check on internal validity, we exogenously provide information about the consequences of both tax regimes and find that such information induces voting for the optimal tax regime. We find suggestive evidence that information provided in the lab on a laboratory referendum affects attitudes on taxation of goods such as electricity or beer.

Our results suggest that underestimation of tax-shifting biases voting, that experience effectively de-biases voters but that deliberation does not improve voting in our laboratory environment. These results are at odds with the standard assumption that voters correctly anticipate the effects of taxing markets and call for taking misperception and learning into account. Below, we motivate some of our design choices and discuss some related literature.

Our first main result is that underestimation of tax-shifting biases individual voting and the choice of tax regimes. Our result is surprising given that the consequences of tax-shifting bias on voting are *a priori* indeterminate. On the one hand, the bias is likely to be common among voters because incentives to exert cognitive effort are inherently weak in voting.

Voters can free-ride on other voters' cognitive efforts because they are unlikely to be pivotal for the outcome. On the other hand, because individual biases aggregate in a non-linear fashion into committee choices, even a considerable share of voters prone to tax-shifting bias may not matter for voting outcomes, as long as biased voters are in minority.

While taxation is often complex in the field³, we believe that the empirical demonstration of a tax-shifting bias is particularly convincing and striking when the environment is simple and relatively transparent. We therefore implement a simple commodity (per unit transaction) tax in a single market. Our demonstration of the existence of biased voting on tax regimes in a controlled environment is novel and we are among the first to study the interaction of voting and markets in the lab (see Grosser and Reuben, 2010 for a notable exception). Yet, we believe that investigating the robustness of our result to de-biasing mechanisms is also important. We investigate the effects of pre-vote deliberation because of its potential to de-bias voters and because of its practical relevance.⁴ We study deliberation in the guise of public, non-structured, free, and non-strategic discussion between voters by relating the content or “quality” of messages to voting.⁵ Deliberation may improve voting outcomes in heterogeneous committees, i.e. when some voters are prone to tax-shifting bias effects while others are not, if the “smart few” persuade at least a majority about the right course of action. In this case, deliberation de-biases committees even if a majority of voters initially falls prey to tax incidence misperception. However, deliberation may also induce “group polarization” which makes initially held opinions more extreme rather than correct. According to Glaeser and Sunstein (2009: 269) there is “pervasive evidence of group polarization on issues that bear directly on politics and political behavior”, and these facts “appear to cast doubt on the wisdom, and certainly the moderation, of crowds”.

3 For example, Kotlikoff and Rapson (2006:1) argue that “thanks to the incredible complexity of the U.S. fiscal system, it's impossible for anyone to understand her incentive to work, save, or contribute to retirement accounts absent highly advanced computer technology and software.” A similar argument is made by Liebman and Zeckhauser (2004) for tax and other (e.g. telephone) schedules.

4 Deliberation is common in many types of committees. For example, trial jurors converse before casting their votes, hiring committees deliberate before making their final decisions, and top management teams hold meetings before determining their firm's investment strategies (see Gerardi and Yariv, 2007 for a discussion). Deliberation in the guise of political discussion is key to the political process. For example, Huckfeldt and Sprague (1987: 1197) note that “politics is a social activity imbedded within structured patterns of social interaction. Political information is conveyed not only through speeches and media reports but also through a variety of informal social mechanisms – political discussions on the job or on the street ... even casual remarks.”

5 There is a considerable literature in experimental economics investigating how the option to communicate affects choices but most studies do not investigate the effects of the content of messages. Notable exceptions are Brandts and Cooper (2007), Charness and Dufwenberg (2006) and Ellingsen and Johannesson (2004).

Our finding that easy-to-interpret information about the consequences of alternative tax regimes changes behavior is noteworthy. While exogenous information has been shown to induce optimal behavior in abstract settings in the laboratory (e.g. Friedman, 1998), evidence that information about fiscal parameters induces optimal behavior is scant. Chetty and Saez (2009) show in a randomized field experiment that providing low-income earners with exogenous information about incentives of a cash transfer program with a relatively complicated incentive structure (EITC) had a marked effect on labor supply and contributed to the reduction of extreme poverty. Information was provided in a personalized and easy-to-interpret way. For example, if marginal incentives conditional on the location in the EITC schedule were appropriate, the information was simply “it pays to work more!”.

Tax-shifting bias as discussed here is related to a literature going back to at least John St. Mill on “fiscal illusion” which discusses misperception of fiscal parameters more generally, including government spending and debt (e.g. Sausgruber and Tyran, 2005; for a survey see Dollery and Worthington, 2006). The tax-shifting bias is also related to but distinct from tax salience. The notions are related because both lead to misperception of the tax burden but they are distinct because tax-shifting bias pertains to taxing markets exclusively while misperception of effective taxes may also occur in individual decision making contexts. Accumulating evidence suggests that misperception of tax incentives due to salience effects is common and can distort choices. Chetty et al. (2009) show in a field experiment that sales taxes that are explicitly stated on price tags in a grocery store (and are thus more salient) have larger effects on demand than if the tax is only added at the cashier. Finkelstein (2009) finds that electronic toll collection is less salient than when cash has to be handed at the toll station, and that reduced salience induces higher toll rates. Salience effects can bias consumers’ perception of the “net” or “final” price and make it less salient than the gross price. This may happen, in the diction of Gabaix and Laibson (2006), when the net price differs from the gross price by some “shrouded attribute” (see Hossain and Morgan, 2006, for a field experiment on the effect of shipping fees in eBay auctions, and DellaVigna, 2009 for a review of this literature). Misperception of taxes, however, may not necessarily be due to salience effects. Blumkin et al. (2008) find that theoretically equivalent consumption and income taxes are behaviorally non-equivalent due to myopic behavior in an individual decision making laboratory experiment. Some evidence suggests misperception of estate tax rates (Slemrod 2006), and there is suggestive evidence of misperception of marginal tax incentives (Liebman and Zeckhauser, 2004 and Feldman and Katuscak, 2007).

2 Experimental Design

In our experiment, participants earn money by trading a hypothetical good in an initial market phase 0 and then vote on imposing a commodity tax to be levied on transactions in the following market phase. Participants are informed that markets need to be taxed and that all tax revenues are wasted, i.e. not returned to participants in any way and that only buyers can vote. In the main treatment BuyerTax, the choice is between a buyer tax of $t = 25$ and a seller tax τ of more than 25. The level of the seller tax is endogenously determined (see section 2.2 for details).

According to standard economic theory, higher taxes translate into lower market rents for consumers in a market with flexible prices (see Fullerton and Metcalf, 2002). This relation holds irrespective of whether the tax is levied on the buyers or the sellers and follows from Tax LSE.

Figure 1: Induced supply and demand: buyer tax (left) and seller tax (right)

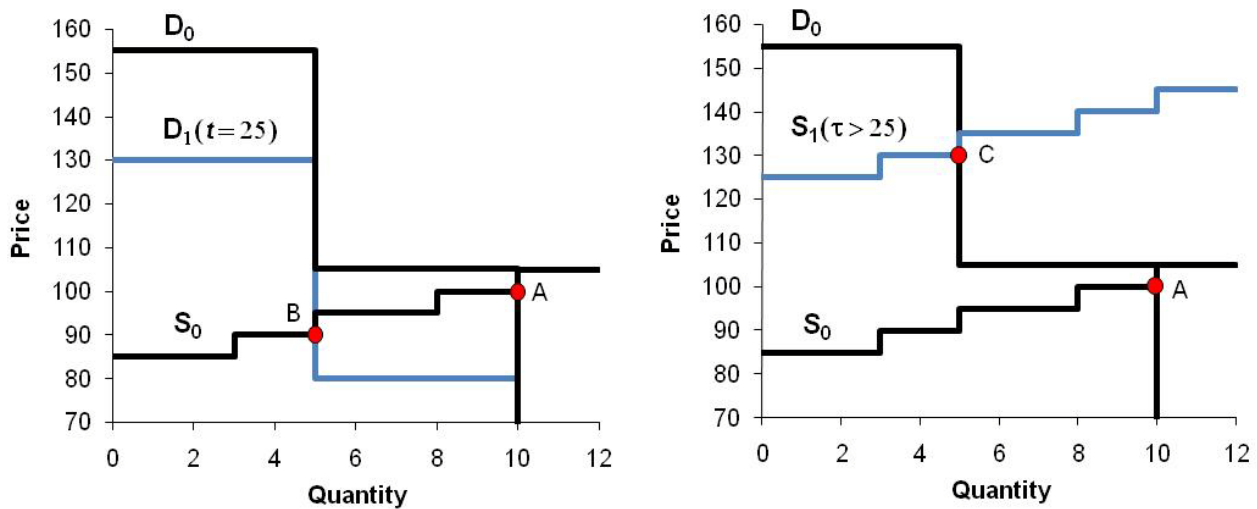


Figure 1 illustrates the consequences of taxation in the two tax regimes. S_0 and D_0 show supply and demand prior to taxation and the pre-tax market equilibrium is at point A, at a quantity of $q_0 = 10$ units and a price of $p_0 = 100$. In the buyer tax regime (see left panel) demand shifts down to D_1 resulting in equilibrium B with $q_t = 5$ and $p_t = 90$.⁶ In the seller tax regime with $\tau > 25$, supply shifts up, resulting in the after-tax equilibrium C with the same quantity of $q_t = 5$ as in case of the buyer tax. However, with $\tau > 25$ prices increase by more

⁶ Note that equilibrium prices fall by less than 25 points due to partial tax shifting. In Sausgruber and Tyran (2005) we investigated the more extreme cases of no or complete tax shifting.

than 25 points and consumer rent is smaller than with the buyer tax of $t = 25$. Hence, in competitive equilibrium $t = 25$ dominates $\tau > 25$ from the consumers' perspective.

Note that voters only need minimal information to correctly predict that their earnings are smaller with the seller tax. For example, voters do not need to know overall market parameters (e.g. demand and supply elasticities) or how much of the tax is shifted.⁷ All they need to know is that higher taxes reduce their incomes and that market prices respond quickly to shifts in demand and supply.

Our experiment has 3 treatments (2 additional treatments are discussed in section 3.4). The baseline treatment BuyerTax is as explained above. Treatment Deliberation is the same as treatment BuyerTax, except that the voting stage is preceded by a deliberation stage and serves to isolate the de-biasing effects of deliberation. Treatment SellerTax is the same as BuyerTax, except that the default tax is reversed. In SellerTax, the choice is between a seller tax of 25 and a buyer tax of more than 25. Treatment SellerTax serves to identify the bias.

To study the de-biasing effects of experience, we repeat the main phase 5 times. Each main phase consists of 2 stages in treatment BuyerTax: A voting stage and a market stage. The voting stage is explained both orally and in writing before the first vote (see Appendix C for instructions). Participants learn the outcome of the ballot (median tax vote and τ) at the end of the voting stage. The market stage has 10 periods with a buyer or a seller tax regime depending on the outcome of the voting stage. Since experience is endogenous in phases 2-5, the causal effect of information about the adverse consequences for voting non-optimal taxes is difficult to isolate. We therefore run two additional treatments that serve, among other things (see section 3.4 for explanations), to isolate the causal effect of such exogenously provided information.

2.1 Details on the market mechanism

We use a uniform-price sealed bid/offer auction (a simplified version of Smith et al., 1982) with the supply and demand parameters shown in Figure 1 in all treatments. Using this market institution has three important advantages for our purposes. First, the market is known from previous experiments to quickly converge to equilibrium predictions implying that tax-shifting bias is costly in our design. Second, trading in our market is simple and is easy to explain to participants, allowing participants to concentrate on the choice of the tax regime

⁷ The extent of tax shifting is market specific. Besley and Rosen (1999) analyze shifting of sales taxes in the US and “find a surprising variety of shifting patterns.” (p. 158). A study for France (Carbonnier, 2007) finds that a specific sales tax is shifted by 57% for new cars, and by 77% for housing repair services.

which is the main focus of our paper. Third, this market institution allows us to automate sellers which simplifies the analysis and enables us to isolate tax-shifting bias salience from fairness considerations.

The details of the market institution are as follows. A market has five buyers and automated sellers. Each buyer has privately known reservation values for two units $v_1 = 155$ and $v_2 = 105$. Buyers submit bids for each unit they want to buy. Buyers can bid at most their induced value per unit. The automated sellers have costs per unit ranging from 85 points to 100 points and submit offers at true cost for each unit to sell. Buyers do not know the unit costs of the sellers but do know that sellers submit offers at the true cost. We automate sellers to control for fairness considerations. Kerschbamer and Kirchsteiger (2000) argue that tax LSE may fail to hold if subjects perceive shifting the tax burden as unfair. Since we automate the sellers and since this is known to buyers, fairness considerations should not affect voting, allowing us to study tax-shifting bias in isolation.

Market outcomes are determined as follows: Bids are ordered from high to low, and offers from low to high. All bids exceeding offers are accepted, and the last accepted bid determines the uniform price for all units traded. A buyer's payoff per traded unit is the difference between the induced value and the market price or zero, if he does not trade. In our experiment, both buyer and seller taxes are transaction taxes, i.e. the tax only depends on the number of transactions, not on their value. We deliberately chose a simple tax structure to make tax incentives transparent. A tax on sellers adds to their cost. Since automated sellers bid their true (tax inclusive) cost, implementing a seller tax of τ shifts the supply function in figure 1 up by τ . Similarly, a tax on buyers shifts the demand function down by t .

Note that the market price is determined by the last accepted bid. Therefore, the bids of the (marginal) buyer affects market prices and buyers experience price fluctuations in the 10 periods of trading in phase 0. Phase 0 serves to familiarize participants with trading in the market and to make sure that participants experience that prices are flexible - which is a necessary condition for tax LSE to hold. Indeed, in phase 0 participants can see that market prices depend on participants' choices and that prices are flexible. In period 1 of phase 0, the variance of market prices is 5.3 points. Phase 0 also sets the stage for our analysis of voting behavior because markets converge well to equilibrium predictions. For example, market efficiency, defined as actual over equilibrium rents, averaged across all treatments is 98.4 percent in the last period of phase 0.

2.2 Details on voting

In treatments BuyerTax and Deliberation, the choice is between a buyer tax of 25 and a seller tax of more than 25. Rather than having voters choose between two exogenous taxes, we let them choose between an exogenously given buyer tax of 25, and an endogenously determined seller tax. The value of the seller tax is determined by an incentive-compatible mechanism designed to elicit the maximum seller tax a buyer is willing to accept to avoid the buyer tax. The purpose of endogenizing the seller tax is to induce voters to think about how much they prefer the seller tax over the buyer tax and to provide them with incentives to state their true preference.

The details of voting are as follows. Voters are in committees of five. Only buyers can vote and voting is compulsory. Voters simultaneously choose integer seller tax rates τ_i between 0 and 50 points.⁸ A random draw τ from a uniform distribution with support $[25, 50]$ determines if the seller tax regime is implemented and at what rate. If the random draw is below the median vote, the seller tax regime is implemented at rate τ . If the random draw is above the median vote, the buyer tax regime is implemented at rate $t = 25$. Therefore, the implemented seller tax is above 25 by design if the seller tax regime prevails in the main treatments. If the choice of the median voter is 25 or less, the committee is said to have voted for the buyer tax regime and the buyers are taxed at $t = 25$ in the following trading phase in this case. If the median choice is above 25, the committee is said to have voted for the seller tax regime and sellers are taxed at the randomly determined rate τ if the median voter was willing to accept a seller tax of at least τ .⁹ Intuitively, the mechanism can be thought of as an incentive-compatible shortcut to majority voting. It is equivalent to the outcome of a vote between t and an exogenous τ .

The implemented seller tax is determined by a random draw to make voting incentive-compatible. Suppose a voter thinks as the standard economist who is indifferent between a buyer tax of 25 and a seller tax of 25. Such a voter casts a vote of 25 or less. Suppose a voter is prone to tax-shifting bias and believes that taxes on buyers fall to a greater extent on buyers than equivalent taxes levied on sellers. Such a voter would be indifferent between the buyer tax of 25 and a seller tax of more than 25. Our mechanism provides incentives to cast a vote

8 Admitting votes for seller taxes below 25 avoids confound with unsystematic decision errors. To see why, assume that we had confined voting choices to the interval $[25, 50]$ and assume that voters state $\tau_i = 25 + \varepsilon$, where ε is a random component following some symmetric distribution. In this case, the median vote would suggest a bias of $\varepsilon/2$ despite errors being unsystematic.

9 To ease the understanding for experimental subjects, the instructions refer to the buyer tax regime as “Party A” and to the seller tax regime as “Party B”. See Appendix D for instructions.

according to his indifference point rather than to strategically shade his vote.¹⁰ The intuition why the mechanism is incentive compatible is that the vote τ_i only affects the probability that the seller tax is implemented at a subjectively preferred rate; it does not affect the distribution from which the seller tax rate is drawn (see Appendix A for a proof).¹¹

The main advantage of this procedure is that it provides incentives to state the true preferences and that it is relatively simple to explain. The drawback is that observed outcomes (i.e. τ) are random and always below the median vote by design. We are therefore careful in distinguishing between individual votes, median (committee) votes and implemented taxes in our discussion of results.

Treatment SellerTax serves to identify the effect of the tax-shifting bias on voting. Treatment SellerTax is exactly the same as BuyerTax except that the default is now a tax on the sellers at $t = 25$ and the vote is on whether to have a tax above 25 on buyers. Now, it is the buyers who are taxed at the rate $\tau > 25$ (again determined by the incentive compatible mechanism) if the median voter was willing to accept a buyer tax of at least τ (see instructions in Appendix D for details). The treatment serves to identify the effect of the tax-shifting bias on voting under the assumption that random voting is equally common in the two treatments. The intuition is that the tax-shifting bias induces consumers to vote for high taxes on sellers but not on buyers. Thus, according to the hypothesis, the bias is effective in treatment BuyerTax but not in SellerTax. If voters vote for taxes above 25 for reasons unrelated to the tax-shifting bias (e.g., unsystematic random errors), subtracting the tax votes in SellerTax from those in BuyerTax gives an estimate of the effect of the tax-shifting bias on voting.

2.3 Details on deliberation

Treatment Deliberation is the same as treatment BuyerTax, except that a deliberation stage precedes voting. Deliberation is organized via computerized, written chat. Participants can send messages which are posted to all members of their committee at any time during 5 minutes. We ask subjects not to reveal their identity, induced values or personal information (participants have fixed ID numbers, see Appendix B), and not to threaten or insult other

10 To illustrate that incentives to vote strategically prevail absent the random draw, suppose a voter is indifferent between a buyer tax of 25 and a seller tax of 35, say. If the voter thinks that the seller tax is partly shifted to the buyers, he has an incentive to vote for the minimal seller tax guaranteeing that the seller tax regime is implemented (i.e. 26).

11 The same intuition explains why an agent has an incentive to bid his true value in the second price auction. There, the agent chooses a bid which maximizes his expected payoff when the price (the second highest bid) is independent of the own bid.

subjects. The experimenter monitors participants' messages online to enforce compliance with these rules.

Deliberation is public, non-structured, free, and non-strategic. It is public in the sense that voters post messages on a (within-committee) public chat board. It is non-structured in the sense that messages are not constrained with respect to origin, sequence, frequency, and length and no priority is given to particular senders or messages. Deliberation is free as there is no cost of sending or receiving messages. Finally, deliberation is non-strategic in the sense that voters are in a common interest situation, i.e. face exactly the same incentives in voting (recall that only buyers can vote).

A total of 285 subjects participated in our experiment as follows: 75 subjects in treatment BuyerTax, 60 subjects in Deliberation, and 70 in the control treatment SellerTax. An additional 40 subjects each participated in two control treatments which serve to test issues of external and internal validity (see section 3.4 for details). Participants were undergraduate students from various majors at the University of Innsbruck. An experimental session lasted approximately 2 hours. The average participant earned €22.6, including a show-up fee of €4. The experiments were run using the software z-Tree (Fischbacher, 2007).

3 Results

Our main results are as follows. Section 3.1 shows that voters are systematically biased absent deliberation and experience. Section 3.2 shows that voters tend to be influenced by what others say, i.e. we find that the quality of messages received is related to the recipient's voting. We find that deliberation tends to make initially prevailing opinions more extreme rather than more correct. Overall, the positive and negative effects of deliberation cancel out, and deliberation does not eliminate tax-shifting bias. Section 3.3 discusses the effects of learning from experience. We find that the tax-shifting bias is rather persistent but that experiencing the adverse effects of high taxes improves voting. However, because voters only learn about the consequences of a tax regime if they chose it in the first place, it is not possible to cleanly isolate causal effects of experience on voting in this setting. Section 3.4 therefore presents results from two additional treatments showing that exogenously provided information about the consequences of alternative tax regimes causally improves voting in the lab. These treatments also provide suggestive evidence on external validity. We find that voting in the lab is related to attitudes on taxation outside the lab and that these attitudes are systematically affected by the exogenous information provided in the lab.

3.1 Tax-shifting bias among inexperienced voters absent deliberation

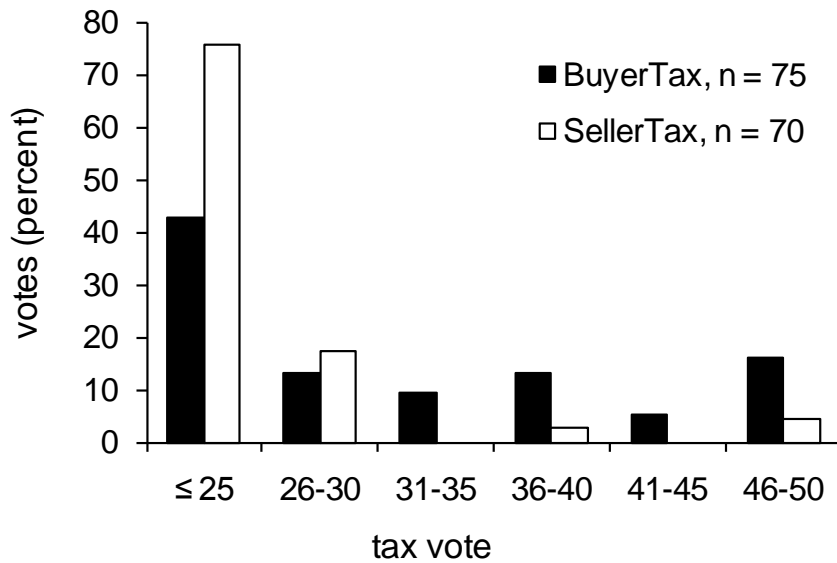
We now show that inexperienced voters were prone to tax-shifting bias absent deliberation. That is, we show that buyers systematically vote for inefficiently high seller taxes in phase 1. We test for the bias by comparing voting in treatment BuyerTax and SellerTax which both had no pre-vote deliberation. The two treatments differ with respect to the default tax. In BuyerTax, the default is a buyer tax of $t = 25$ and participants vote on a seller tax. In SellerTax, the default is a seller tax of 25 and participants vote on a buyer tax.

Standard economics predicts that voters prefer low taxes, i.e. do not vote for taxes above 25 in both treatments. In contrast, the tax-shifting bias hypothesis predicts an asymmetric distribution of votes in the two treatments. According to this hypothesis, voters prefer to avoid the buyer tax in both cases. In particular, voters are predicted to vote for high (seller) taxes to avoid the buyer tax in BuyerTax, but to vote for low (buyer) taxes to avoid the buyer tax in SellerTax. Thus, the tax-shifting bias hypothesis predicts a higher share of votes above 25 in BuyerTax than in SellerTax.

Figure 2 shows the distribution of tax votes of inexperienced voters in treatments BuyerTax and SellerTax. The figure shows that optimal voting is more pronounced in SellerTax (76 vs. 43 percent), and that, conversely, tax votes above 25 are more than twice as common in BuyerTax than in SellerTax (57 vs. 24 percent). The average tax vote¹² is significantly different (33.1 vs. 26.8, $p = 0.000$, Mann-Whitney (MW) test) and the distribution of tax choices is significantly different across treatments ($p = 0.002$, Kolmogorov-Smirnov (KS) test). We conclude that voters are heterogeneous with respect to how prone they are to tax-shifting bias and there is a fair amount of non-optimal, perhaps random, voting (24 percent) even in SellerTax. If we assume the same incidence of non-optimal voting in BuyerTax that comes from other sources than tax-shifting bias, we find that about a third (= 57% - 24%) of non-optimal tax votes in BuyerTax can be attributed to systematic tax-shifting bias.

12 In calculating average tax vote we count all tax votes below 25 as equal to 25. The reason is that stating a tax vote below 25 rather than exactly 25 does not affect which tax regime is implemented.

Figure 2: Distribution of tax votes of inexperienced voters (Phase 1)



We find that committee choices are much more in line with standard predictions in SellerTax than in BuyerTax, as predicted by the tax shifting hypothesis. In SellerTax, only 1 out of 14 committees (7.1 percent) has a median tax vote > 25 and the median tax vote averaged across all committees is 25.1 points, i.e. very close to the rational choice benchmark of 25. In contrast, in BuyerTax, two thirds (10 out of 15) of all committees have a median tax vote > 25 and the median vote averaged over all markets is 31.5 points. The share of committees accepting taxes above 25 is significantly higher in BuyerTax than SellerTax ($p = 0.001$, χ^2 test), and the median tax votes are significantly higher in BuyerTax than SellerTax ($p = 0.000$, MW).

Committee votes are linked to market outcomes through the incentive-compatible voting mechanism which involves a random draw to induce voters to reveal their true preference. While there was one committee that voted for taxes above 25 in SellerTax, no committee experienced such a tax in phase 1. In BuyerTax, 10 committees voted for taxes above 25 and 3 markets experienced such a tax (with an average tax rate of 31.3).

How do committee votes translate into market outcomes? As predicted by standard theory, higher taxes reduce market incomes. In particular, buyers in markets with taxes above 25 in BuyerTax earn less than with taxes of 25 in BuyerTax (25.9 vs. 32.6, $p = 0.061$, MW) and in SellerTax (25.9 vs. 32.2, $p = 0.058$, MW). Also in line with standard theory, we find that buyer and seller taxes at a given rate have the same market consequences, i.e. tax LSE holds. In particular, we find that buyer earnings in markets with a tax of 25 do not differ

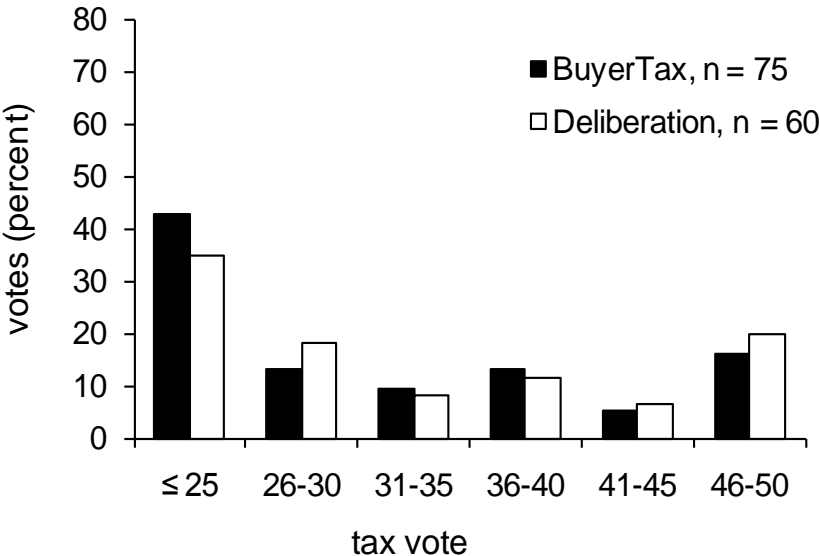
between BuyerTax and SellerTax (32.6 vs. 32.2, $p = 0.897$, MW). We conclude that our market institution generates outcomes predicted for perfectly competitive markets, a finding that is in line with results from numerous experiments (see footnote 2 for references). For the sake of brevity, we therefore abstain from providing a detailed discussion of how taxation affects prices and quantities and focus on the effects of deliberation instead.

3.2. Deliberation and tax-shifting bias

Does pre-vote deliberation mitigate the tax-shifting bias? To investigate, we compare treatments BuyerTax and Deliberation for inexperienced voters. Treatment Deliberation is the same as BuyerTax, except that voters are given the option to communicate during 5 minutes by sending each other messages before the vote. Recall that deliberation is public, non-structured, free, and non-strategic. Deliberation is also voluntary in the sense that voters cannot be forced to send messages or to pay attention to messages they receive. Yet, voters used the option to deliberate extensively. In the deliberation stage of phase 1, the average number of messages sent per voter is 6.5, and the median is 5.0.

Figure 3 shows that pre-vote deliberation did not affect the distribution of individual tax votes. In Deliberation, 65 percent of tax votes are above 25, and the average tax vote is 34.1 points. In BuyerTax, 57 percent of tax votes are above 25 and the average tax vote is 31.5. Neither the average tax vote ($p = 0.696$, MW) nor the distribution of tax votes above 25 is significantly different in Deliberation and BuyerTax ($p = 0.839$, KS).

Figure 3: Distribution of tax votes with and without deliberation (Phase 1)



How are committee votes affected by the option to deliberate? The short answer is that they are not because the individual-level effects of deliberation (see below) tend to cancel out. In BuyerTax, two thirds of all markets (10 out of 15) have median tax votes above 25, and the median tax vote averaged across all markets is 31.5 points. The results for Deliberation are very similar. In Deliberation, again two thirds of markets (8 out of 12) have a median tax vote above 25, and the median tax vote averaged across all markets is 33.3 points ($p = 0.709$, MW).

Deliberation and persuasion

While the conclusion from the previous paragraph is that the option to deliberate had no effect on voting and market outcomes, it is not the case that deliberation had no effect at all. We now show, by analyzing the content or “quality” of messages, that deliberation was persuasive but that correct messages were not more persuasive than incorrect ones, inducing group polarization. Apparently, positive and negative effects of pre-vote deliberation happened to cancel out.

To investigate the effect of the content or “quality” of messages on voting and market outcomes, we classify messages as follows. We call a message “related” if it directly refers to the tax proposal and “unrelated” otherwise. We find that 90.4 percent of all messages were “related” in phase 1 indicating that messages were not just casual conversation but involved problem-oriented deliberation. We call a related message “correct” if it is compatible with the standard economics prediction or proposes an improvement of the committee decision, and “incorrect” otherwise. Finally, we call a related message “none” if it is neither correct nor incorrect (see Appendix B for an example of our classification).

In all committees in Deliberation, at least one correct message is sent. Therefore, if there is something like a “heureka effect” according to which some voters are inattentive but recognize the optimal action if alerted to it, we would expect all committees in Deliberation to vote in line with tax LSE. As shown later in more detail, this is not the case. We argue that this is the case because in all committees voters are exposed to some correct and some incorrect messages in varying proportions. In fact, the average shares of correct and incorrect messages per committee are not significantly different according to a Wilcoxon signed-rank test (22.3 vs. 32.5 percent, $p = 0.308$).

How does the quality of messages correlate with the quality of committee decisions? We find that these variables are strongly related. The more correct messages are exchanged in a committee, the better the committee decision. In fact, the larger the number of correct

messages, the lower the committee's tax vote (Spearman's rho = -0.673, $p = 0.016$), and vice versa for incorrect messages (Spearman's rho = 0.856, $p = 0.000$). While these correlations are impressively strong, they need to be interpreted with care. On the one hand, correct messages may tend to reduce, and incorrect messages tend to enforce tax-shifting bias through persuasion. On the other hand, voters who are prone to the bias may also be prone to sending more incorrect messages and vice versa. In this case, the correlation should not be interpreted as a causal effect of deliberation on tax choices but as a selection effect.

As we show next, both types of effects are likely to be present. We distinguish between "persistent senders" who send either correct or incorrect messages throughout, and "fickle senders" who change their opinion over the course of deliberation. We find that "persistent senders" tend to vote in line with what they say but that "fickle senders" tend to vote in line with what other voters say. To demonstrate that this is the case, we group persistent senders into "correct senders" (12/60 subjects send at least one correct and no incorrect message) and analogously defined "incorrect senders" (15/60 subjects), respectively. 87 percent of incorrect senders vote for taxes above 25, but only 17 percent of the correct senders do. In contrast, fickle senders' votes are affected by what consistent senders say.

The presence of correct senders improves voting of fickle senders. While 90 percent (= 18/20) of fickle senders in a committee without any correct sender vote for taxes above 25, only 46 percent (= 6/13) do so in the presence of at least one correct sender. In contrast, the presence of incorrect senders tends to deteriorate voting of fickle senders (60% vote for taxes above 25 absent, 78% in the presence of at least one incorrect sender). Thus, fickle senders are persuaded by messages from both correct and incorrect senders. These results are confirmed by a Probit regression (using all subjects) with voting for taxes above 25 as the dependent variable. We find that (controlling for the number of messages exchanged) a one-percent increase in the share of correct messages sent by others within the committee reduces the probability to vote for taxes above 25 by 0.29 percent ($p = 0.006$) while a one-percent increase in the share of incorrect messages increases this probability by 0.15 percent ($p = 0.025$).

Deliberation and group polarization

The analysis above suggests that messages are persuasive, i.e. that voting is systematically affected by the quality of messages received. We show next that deliberation induces voters to agree – but not necessarily on the right solution. Instead, deliberation tends to make initially held opinions more homogenous and more extreme – a process called group

polarization (e.g. Isenberg, 1986). Committees that happen to exchange many incorrect messages tend to vote for taxes above 25 and vice versa for correct messages. As a result, we find that the option of pre-vote deliberation does not mitigate the tax-shifting bias among inexperienced voters.

Agreement is measured by the within-committee standard deviation of individual tax votes. The average within-committee standard deviation is almost twice as large in BuyerTax as in Deliberation (14.8 vs. 7.9, $p = 0.007$, MW). We group committees into low and high agreement depending on whether a committee has a standard deviation above or below the median committee in BuyerTax and Deliberation. According to this definition, 83 percent of committees have a high level of agreement in Deliberation, but only 27 percent of the committees in BuyerTax do ($p = 0.005$, χ^2 -test). While deliberation induces voters to agree, they do not systematically agree on the lower tax. In fact, the correlation between committee votes and within-committee standard deviation is insignificant (Spearman's rho = -0.318, $p = 0.313$).

While deliberation reduces the within-committee standard deviation of taxes votes, it increases its across-committee standard deviation. It is more than twice as large in Deliberation (11.8) than in BuyerTax (5.8). If we consider a session an independent observation (we had 3 sessions per treatment), we find that across-committee standard deviations were significantly higher in Deliberation than in BuyerTax ($p = 0.063$, one-sided MW). Thus, deliberation seems to induce a “group polarization” effect.

3.3 Learning, deliberation and tax-shifting bias

While the previous sections focused on the behavior of inexperienced voters (in phase 1), we now discuss behavior in later phases. However, the results reported below should be interpreted with due care due to the endogeneity of experience. For example, mean reversion is likely to occur because committees making particularly bad choices in one phase are more likely to learn the adverse effects of high taxes than those voting for moderate taxes. To unambiguously identify the causal effect of information about the adverse effects of high taxes, we discuss the effects of exogenous information in section 3.4.

Our main findings are that the tax-shifting bias is surprisingly persistent, that learning is promoted by “informative experience” (explained below), and that deliberation has no significant effect on voting behavior throughout all phases. To illustrate the last point, we note that voting is not significantly different with and without deliberation both at the individual and the committee level over all phases. In particular, the average vote over all phases is 33.1

in Deliberation vs. 32.8 in BuyerTax ($p = 0.941$, MW). The average committee vote is 32.0 points in Deliberation vs. 30.8 points in BuyerTax ($p = 0.526$, MW). We find some evidence of convergence to optimal voting but convergence is slow in both BuyerTax and Deliberation. For example, the share of individual votes for taxes above 25 falls in Deliberation from 65 to 42 percent and in BuyerTax from 57 to 40 percent from phase 1 to 5. The share of groups with committee votes above 25 falls from 67 to 42 percent in Deliberation and from 67 to 52 percent in BuyerTax. In both treatments, the share of markets experiencing the seller tax regime falls by about a third (from 25 to 17 percent in Deliberation and from 25 to 13 percent in BuyerTax) from phase 1 to 5.

The ineffectiveness of deliberation in later phases of the experiment is not due to participants simply stopping to deliberate or messages becoming noisy. For example, the average number of messages a subject sends is not significantly different in phase 5 from phase 1 (5.7 vs. 6.7, $p = 0.530$, Wilcoxon signed-rank test). As in phase 1, the quality of messages exchanged correlates well with the quality of committee decisions in the remaining phases. We find that the more correct messages are exchanged, the better the committee choice averaged over phases 2 to 5 (Spearman's rho = -0.724, $p = 0.007$), while the reverse is true for incorrect messages (Spearman's rho = 0.718, $p = 0.008$).

Given the similarity of dynamics in BuyerTax and Deliberation, we merge the data from BuyerTax and Deliberation and compare it to SellerTax to test for the persistence of tax-shifting bias. Recall that according to the bias hypothesis, BuyerTax and Deliberation induce voting for high taxes but SellerTax does not. Over all phases, BuyerTax and Deliberation have significantly higher committee votes (BuyerTax and Deliberation: 31.3, SellerTax: 26.7, $p = 0.001$, MW), higher tax revenues (BuyerTax and Deliberation: 127.5, SellerTax: 121.3, $p = 0.041$, MW), and lower efficiency (BuyerTax and Deliberation: 94.4%, SellerTax: 95.8%, $p = 0.099$, MW) than SellerTax.

Informative experience prevails if the committee vote was for a tax above 25 and the committee experienced a tax above 25 in the market stage in a previous phase. In this case, voters experience losses in the market stage. Our results show that informative experience does reduce tax-shifting bias. Almost all committees make an informative experience at some point during the experiment. In particular, only one of the 27 committees in BuyerTax and Deliberation never experiences taxes above 25. However, most markets experience a tax above 25 only once (67 percent in both NO and Deliberation), indicating that they managed to avoid the inefficient tax after having experienced its adverse effects. In phases 2-5, there are

33 instances of informative experience in BuyerTax and Deliberation. In 94 percent (= 31/33) of these cases, the committee vote responds in the right direction, i.e. by voting for lower taxes in the following phase. This finding not only shows that voters managed to respond rationally in the face of easy-to-interpret evidence but also that our experiment provided sufficient incentives for learning given informative feedback.

Finally, we find evidence that experience affects the quality of messages in Deliberation. Without informative experience, the share of correct messages falls and the share of incorrect messages tends to increase, while with informative experience, the share of correct messages increases and the share of incorrect messages falls.¹³ In later phases of the experiment committee polarization effects cede to exist. This is not surprising as subjects become increasingly experienced. For instance, in phase 5, average within-committee standard deviation is 4.5 points in BuyerTax and 6.1 points in Deliberation ($p = 0.237$).

3.4 Testing for internal and external validity

Concerning internal validity, this section shows that providing exogenous information about the consequences of alternative tax regimes reduces the tax-shifting bias in the laboratory. Concerning external validity, we show that voting choices in the laboratory are correlated to attitudes on taxation “in the wild”, and that providing exogenous information about the consequences of alternative tax regimes in the laboratory does affect attitudes on tax issues “in the wild”. This evidence suggests that participants think that the controlled mechanisms applying in the lab may also apply “in the wild”.

To address these issues of internal and external validity, we conducted two additional treatments with 40 participants each. The treatments differ exclusively by whether exogenous information about the consequences of alternative tax regimes is provided (see phase 3 in figure 4). The treatment with exogenous information is called Info, the treatment without information NoInfo.

In phase 1, participants are presented with 6 scenarios involving taxation of goods “in the wild” (beer, insurance policies and electricity), asking them directly about their preferences for taxing goods. Note that responses to these scenarios are not incentivized. In each scenario, the choice is between a tax at rate b to be paid by the buyers and a tax at rate s

13 For example, a comparison of the first and the second vote reveals significant effects in a difference-in-difference test. Without informative experience, the share of correct messages falls by 18.2 percentage points, but increases by 20.8 percentage points with informative experience ($p = 0.021$, MW). For the share of incorrect messages, the differences are -1.0 and -43.6 percentage points ($p = 0.013$, MW). However, experience does not affect the absolute number of messages sent.

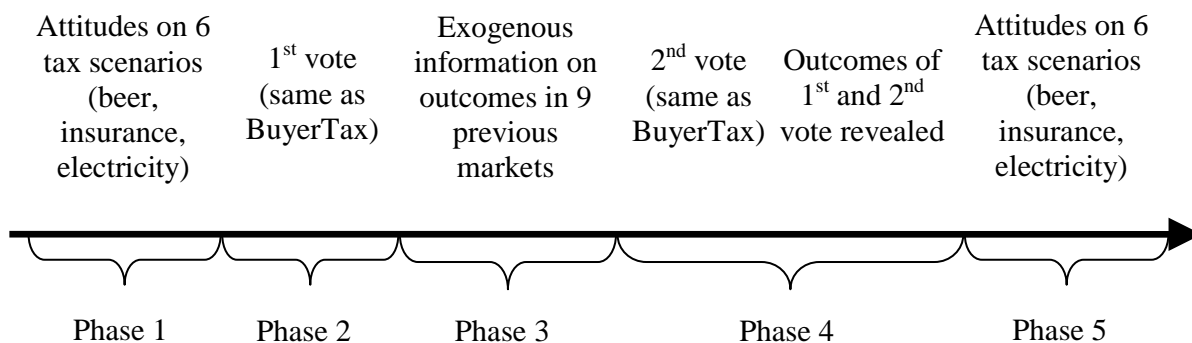
to be paid by the sellers of a particular good. We present respondents with two scenarios for each good in random order. One scenario has $b_1 < s_1$ and the other $b_2 > s_2$, with $s_1 - b_1 = b_2 - s_2$ (see Appendix E).

Phase 2 essentially replicates treatment BuyerTax. That is, participants earn incomes by trading in the market for 10 periods and then vote on the tax regimes, i.e. the choice is between a buyer tax of $t = 25$ and a seller tax of $\tau > 25$. The only difference in procedures to BuyerTax concerns trading after the vote. In contrast to BuyerTax, participants did not trade in the market after the vote in the chosen tax regime in both Info and NoInfo, but earnings were calculated using equilibrium predictions. We take trading out of the hands of participants to keep the experiment short and to remove any strategic consideration which may have affected voting. Participants do not receive feedback about the outcomes of phase 2 (i.e. about the outcome of the vote, market prices and the resulting market income) immediately but only at the end of phase 4 to avoid confounding endogenous experience and exogenous information.

Phase 3 differs across treatments. In Info, voters are provided with exogenous information about the market and payoff consequences of 9 alternative tax choices in previous experiments for both tax regimes. In particular, we sample 9 medians from actual median choices in BuyerTax for inexperienced voters (see Appendix F for details). Since the actual sample has 10/15 markets with seller tax > 25 , this results in two thirds of the cases where the median is > 25 . For each median, we apply the mechanism described in section 2.2 to determine whether the buyer or the seller tax is implemented. Application of this procedure results in about half ($= 4/9$) of the cases with a seller tax regime. These cases are informative in the sense that voters can learn about the adverse consequences of implementing high taxes. For each of the 9 cases, participants are informed about equilibrium market prices and payoffs (we add some noise to make the distribution of market prices and payoffs akin to observed outcomes). In treatment NoInfo, phase 3 is omitted.

Phase 4 is a repetition of phase 2. At the end of phase 4, participants are informed about market outcomes for both phases 2 and 4. Phase 5 is a repetition of phase 1 (recall that scenarios are presented in random order in both phases).

Figure 4: Time line in treatment Info (phase 3 is omitted in treatment NoInfo)



The results from phase 1 show that 58 percent of the responses are in line with tax LSE, i.e. respondents consistently prefer either high or low taxes across a pair of scenarios for a given good, and the vast majority (92 percent) of these consistently prefer the low tax. Among the responses violating tax LSE, i.e. preferring the high tax in one scenario but the low in the other for a given good, we observe a significant bias towards taxing sellers (60 percent of the cases, $p = 0.023$, binomial test). Another way to measure tax-shifting bias is to test if seller taxes are more popular than buyer taxes, given that they are both low. We indeed find that 79% (= 189/240) of responses are in support of seller taxes when $s_2 < b_2$, while support for taxing buyers is weaker (70% = 168/240) when $b_1 < s_1$ ($p = 0.028$, χ^2 test). Thus, violation of tax LSE on tax issues “in the wild” is common and biased towards taxing sellers.

Are attitudes on taxation “in the wild” related to voting choices in the experiment? We indeed find that those who say they prefer low buyer taxes on goods such as beer, electricity or insurance in phase 1 also tend to vote for low buyer taxes in phase 2 (Spearman's rho = 0.091, $p = 0.045$). To provide further support for this claim, we regress voting for the buyer tax in phase 2 on how often a subject expressed a preference for low buyer taxes in phase 1. This regression reveals that the average voter is 14.4 percentage points more likely to vote for the buyer tax in the experiment per additional expression of support for low buyer taxes “in the wild”. That is, the probit regression yields a significant marginal effect of 14.4 ($p = 0.021$) percentage points.¹⁴

Does exogenous information about the consequences of alternative tax regimes in the lab promote optimal voting in the lab? We measure the causal effect of information by testing if optimal voting improves more when information is provided than when it is not, i.e. in a

¹⁴ The regression includes a constant and a dummy for treatment Info. Standard errors are adjusted for 16 clusters within groups, Pseudo $R^2 = 0.08$, Wald χ^2 (df = 2) = 9.13***, $n = 80$.

difference-in-difference approach. In Info, the average change in tax choices from the 1st to the 2nd vote is -5.2 tax points and in NoInfo it is -1.6 points, indicating a significant response of optimal voting to exogenous information ($p = 0.003$, MW one-sided). The same result obtains if we instead calculate the difference-in-difference for the median tax vote (-7.6 vs. -2.0. $p = 0.030$, MW one-sided). The de-biasing effect of information is remarkably strong. The difference between the median tax choices in treatments BuyerTax and SellerTax was $31.5 - 25.1 = 6.4$ points. Now we observe a difference-in-difference effect of learning of $7.6 - 2.3 = 5.3$ points. In other words, exogenous information reduces the tax-shifting bias by 83 percent ($= 5.3 / 6.4$).

Figure 5: Effects of exogenous information on tax-shifting bias

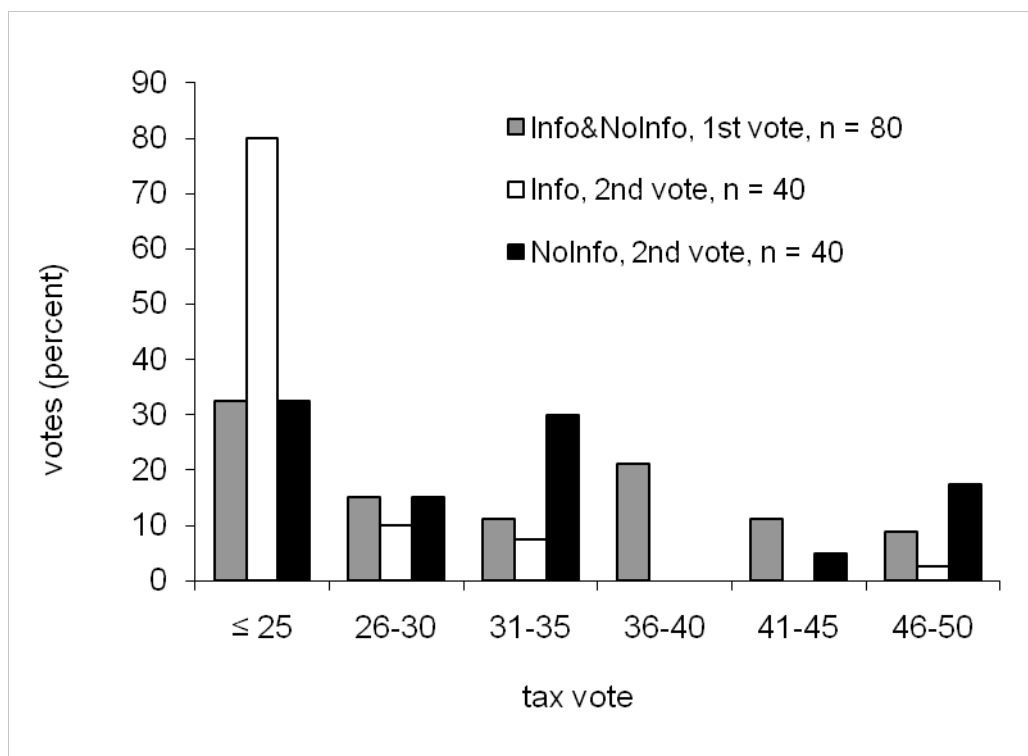


Figure 5 illustrates the strong effect of exogenous information on optimal voting. The figure shows the distribution of votes in Info and NoInfo. The grey bars show the distribution of votes merged for the two treatments when no exogenous information is available in either treatment. This distribution is not different from voting behavior of uninformed voters in treatment in BuyerTax ($p = 0.601$, KS), indicating that the results in Info and NoInfo replicate treatment BuyerTax. Exogenous information had a marked effect in the second vote in Info. In particular, the percentage of optimal votes is 80 percent in treatment Info (see leftmost white bar in figure 5), while it remains at the level of the first vote in treatment NoInfo (see leftmost black bar). Thus, we find that voters strongly respond to exogenous information

about the adverse effects of high taxes but do not seem to recognize the optimal vote if a fellow voter points out the possibility of such effects in a pre-vote conversation (see our discussion of a “heureka effect” in section 3.2). The differential response may be due to the credibility of the information itself (the exogenous information is “hard” empirical evidence rather than a hunch, and it is easy to interpret), due to the clarity of information (exogenous information unambiguously indicates that higher taxes are costly while messages received during deliberation were partly correct and partly misleading), or due to the credibility of the source of information (the experimenter rather than an anonymous fellow voter).

Figure 6: Change in stated preference for the low buyer tax (tax scenarios $b_1 < s_1$)




Figure 6 illustrates “transfer learning” which occurs if subjects take what has been learned in one economic environment and generalize it to related environments (e.g. Cooper and Kagel, 2003). The idea is to test if the provision of exogenous information (in phase 3) about the consequences of taxing laboratory markets also changes attitudes on taxation “in the wild” by comparing the change in individual tax attitudes in Info and NoInfo. To directly address the de-biasing effect of exogenous information, we count the number of times a subject prefers the low buyer tax in the three scenarios where $b_1 < s_1$. Figure 6 shows how these counts change from phase 1 to 5 in Info and NoInfo, where positive numbers indicate that individual attitudes become more consistent with the standard prediction, and vice versa for negative numbers. The figure shows that tax attitudes are generally quite persistent between phases, but persistence is weaker when information is provided (52.5% vs. 67.5% in Info and NoInfo, respectively). Importantly, the figure illustrates that exogenous information induces a more pronounced shift towards a preference for low buyer taxes in Info than in NoInfo ($p = 0.005$, MW one-sided). The result is similar if we instead ask whether exogenous

information induces a general preference for low taxes in all 6 scenarios (rather than low buyer taxes as above). Again, we find that information induces respondents to change their minds since persistence is weaker in Info (37.5% vs. 60% in Info and NoInfo, respectively), and exogenous information promotes a preference for low taxes “in the wild” ($p = 0.050$, MW one-sided).

The observed correlation between attitudes on taxation in the wild and voting behavior in the lab (including the evidence on transfer learning) is remarkable. By virtue of experimental control, the optimal choice is clear in the experiment – participants should vote for lower taxes whether they are levied on buyers or sellers. However, this is not necessarily the case in the field because tax revenues are not (entirely) wasted. A preference for high taxes might well be rational if the respondent does (or believes to) benefit from government spending in some way. Consider, for example, a respondent who does not drink beer but consumes local public goods (higher education, say) funded by the general budget, including beer taxes. However, irrespective of such preferences (or beliefs), rational responses are in line with tax LSE. That is, a respondent who prefers high taxes on some good should be in favor of both high seller and buyer taxes. Yet, this is only the case for a small majority of respondents (58%), and the inconsistent responses tend to be biased towards taxing sellers.

4 Concluding remarks

This paper has shown that tax-shifting bias distorts voting on tax regimes, even if the bias is costly to voters. In our experimental design, voting for the high seller tax to avoid taxing themselves is clearly inefficient because tax revenues are simply wasted. The reason for this design choice is that we wanted to create a tension between tax-shifting bias and efficiency. This tension makes the demonstration of a bias more compelling because voters have an incentive to overcome it. However, whether the bias generally creates inefficiency in the field must remain unclear because tax revenues are not (entirely) wasted and because other biases may be present in the field. On the one hand, voters may have a general aversion to be taxed and underestimate the true (social) benefits of government spending. If so, tax-shifting bias may help to fund socially beneficial expenditure. On the other hand, the tax-shifting bias could allow budget-maximizing politicians and bureaucrats to increase tax revenues beyond what citizens rationally deem appropriate. Some authors have suggested that policy makers may deliberately choose taxes with “hidden” consequences in one way or another to increase tax revenues (e.g. Slemrod and Krishna, 2003; McCaffery and Baron, 2006 and Finkelstein,

2007). However, the extent to which raising revenue is feasible by inducing the tax-shifting bias, i.e. by manipulating voters, depends on political competition (Blumkin and Menirav, 2009).

We believe that accounting for biased perceptions of taxation in the market place is important for understanding the acceptance of alternative tax regimes. We have shown that tax-shifting bias has strong effects in a controlled laboratory setting and provide suggestive evidence that our findings have some external validity. We therefore hope our contribution motivates further investigations into the emerging field of behavioral public economics.

References

- Besley, T., Rosen, H.S., 1999. Sales tax and prices: an empirical analysis. *National Tax Journal* 52(2), 157-78.
- Blinder, A.S., 1988. The challenge of high unemployment. *American Economic Review* 78, 1-15.
- Blumkin, T., Menirav, E., 2009. Framing the rabbit to snare the votes. *Social Choice and Welfare* 32, 603-34.
- Blumkin, T., Ruffle, B.J., Ganun, Y., 2008. Are income and consumption taxes ever really equivalent? Evidence from a real-effort experiment, University of Munich Working paper 6479.
- Borck, R., Engelmann, D., Müller, W., Normann, H.-T., 2002. Tax liability side equivalence in an experimental posted offer market. *Southern Economic Journal* 68(3), 672-82.
- Brandts, J., Cooper, D.J., 2007. It's what you say, not what you pay: an experimental study of manager-employee relationships in overcoming coordination failure. *Journal of the European Economic Association* 5(6), 1223-68.
- Carbonnier, C., 2007. Who pays sales taxes? Evidence from French VAT reforms, 1987-99. *Journal of Public Economics* 91(6), 1219-29.
- Charness, G., Dufwenberg, M., 2006. Promises and partnerships. *Econometrica* 74, 1579-601.
- Chetty, R., Looney, A. Kroft, K., 2009. Salience and taxation: theory and evidence. *American Economic Review* 99(4), 1145-77.
- Chetty, R., Saez, E., 2009. Teaching the tax code: earnings responses to an experiment with EITC recipients. NBER working paper no. 14836.
- Cooper, D.J., Kagel, J.H., 2003. Lessons learned: generalizing across games. *American Economic Review* 93, 202-7.
- DellaVigna, S., 2009: Psychology and economics: evidence from the field. *Journal of Economic Literature* 47, 315-72.
- Dollery, B.E., Worthington, A.C., 2006. The empirical analysis of fiscal illusion. *Journal of Economic Surveys* 10(3), 261-97.
- Ellingsen, T., Johannesson, M., 2004. Promises, threats and fairness. *Economic Journal* 114, 397-420.
- Falk, A., Heckman, J.J., 2009. Lab experiments are a major source of knowledge in the social sciences. *Science* 326(5952), 535-8.
- Feldman, N.E., Katuscak, P., 2007. Effects of predictable tax liability variation on household labor income. Ben Gurion University Working paper.
- Finkelstein, A., 2009: E-ZTax. Tax salience and tax rates. *Quarterly Journal of Economics* 124(3), 969-1010.
- Fischbacher, U., 2007. z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics* 10(2), 171-8.
- Friedman, D., 1998. Monty Hall's three doors: construction and deconstruction of a choice anomaly. *American Economic Review* 88(4), 933-46.

- Fullerton, D., Metcalf, G.E., 2002. Tax incidence, in: Auerbach, A.J., Feldstein, M. (Eds.), *Handbook of Public Economics* Vol. 4, Elsevier, Amsterdam, 1787-872.
- Gabaix, X., Laibson, D., 2006. Shrouded attributes and information suppression in competitive markets. *Quarterly Journal of Economics* 121, 505-40.
- Gerardi, D., Yariy, L., 2007. Deliberative voting. *Journal of Economic Theory* 134, 317-38.
- Glaeser, E.L., Sunstein, C.R., 2009. Extremism and social learning. *Journal of Legal Analysis* 1(1), 263-324.
- Graetz, M.J., 2002. 100 million unnecessary returns: a fresh start for the U.S. tax system. *Yale Law Journal* 112, 261-310.
- Grosser, J., Reuben, E., 2010. Redistributive politics and market efficiency: an experimental study. IZA Working paper 4549.
- Hossain, T., Morgan, J., 2006. ... Plus shipping and handling: revenue (non) equivalence in field experiments on eBay. *Advances in Economic Analysis and Policy* 6(2), Article 3.
- Huckfeldt, R., Sprague, J., 1987. Networks in context: the social flow of political information. *American Political Science Review* 81(4), 1197-216.
- Isenberg, D.J., 1986. Group polarization. *Journal of Personality and Social Psychology* 50(6), 1141-51.
- Kachelmeier, S., Limberg, S., Schadewald, M., 1994. Experimental evidence of market reactions to new consumption taxes. *Contemporary Accounting Research* 10(2), 505-46.
- Kerschbamer, R., Kirchsteiger, G., 2000. Theoretically robust but empirically invalid? An experimental investigation into tax equivalence. *Economic Theory* 16(3), 719-34.
- Kotlikoff, L.J., Rapson, D., 2007. Does it pay, at the margin, to work and save? – Measuring effective marginal taxes on Americans' labor supply and saving. NBER working paper 12533.
- Liebman, J.B., Zeckhauser, R.J., 2004. Schmeduling. Harvard University working paper.
- McCaffery, E.J., Baron, J., 2006: Thinking about tax. *Psychology, Public Policy, and Law* 12(1), 106–35.
- Riedl, A., Tyran, J.-R., 2005. Tax liability side equivalence in gift-exchange labor markets. *Journal of Public Economics* 89(11), 2369-82.
- Ruffle, B., 2004. Tax and subsidy incidence equivalence theories: experimental evidence from competitive markets. *Journal of Public Economics* 89(8), 1519-42.
- Sausgruber, R., Tyran, J.-R., 2005. Testing the Mill hypothesis of fiscal illusion. *Public Choice* 122, 39-68.
- Slemrod, J., 2006. The role of policy misconceptions in support for regressive tax reform. *National Tax Journal* 59(1), 57-75.
- Slemrod, J., Krishna, A., 2003. Behavioral public finance: tax design as price presentation. *International Tax and Public Finance* 10(2), 189-203.
- Smith, V.L., Williams, A.W., Bratton, W.K., Vannoni, M.G., 1982: Competitive market institutions: double auction versus sealed bid-offer auctions. *American Economic Review* 72(3), 58-77.

Appendices

Appendix A: Proof of the incentive-compatibility of the voting mechanism

Assume that a buyer maximizes expected utility and attaches a positive probability to being decisive. A transaction tax τ_i with $i = B, S$ shall be imposed. The index B indicates that the tax is levied on the buyers, index S denotes that the tax is levied on the sellers. The rate of the buyer tax is a constant, $\tau_B = \bar{\tau}_B$. The rate of the seller tax is $\tau_S, \tau_S \in [\bar{\tau}_B, \tau_{MAX}]$. Let $U_B(\tau_i)$ be the buyer's net benefit under tax τ_i . Call $\bar{\tau}_S$ the rate of the seller tax that yields the same net benefit for the buyer as the buyer tax at rate $\bar{\tau}_B$. Assume that

$$\begin{aligned} U_B(\tau_S) &> U_B(\bar{\tau}_B) \quad \forall \tau_S \in [\bar{\tau}_B, \bar{\tau}_S[, \\ U_B(\tau_S) &= U_B(\bar{\tau}_B) \quad \forall \tau_S = \bar{\tau}_S, \text{ and} \\ U_B(\tau_S) &< U_B(\bar{\tau}_B) \quad \forall \tau_S \in]\bar{\tau}_S, \tau_{MAX}]. \end{aligned}$$

The tax is determined by the following mechanism: The buyer announces a rate for the seller tax τ'_S . A uniformly distributed random variable z is drawn from the interval $[\bar{\tau}_B, \tau_{MAX}]$. The implemented tax is

$$\begin{aligned} \tau_B &\quad \forall z > \tau'_S \\ \tau'_S &= z \quad \forall z \leq \tau'_S \end{aligned}$$

Claim: it is optimal for the buyer to announce $\tau'_S = \bar{\tau}_S$.

Proof: With every choice of τ'_S the buyer determines a lottery with two prizes. One prize is $U_B(\bar{\tau}_B)$, the other $U_B(\tau'_S)$. The lotteries fall into two classes. The first is with $\tau'_S \leq \bar{\tau}_S$ and has $U_B(\tau'_S) \geq U_B(\bar{\tau}_B)$. As long as $\tau'_S \leq \bar{\tau}_S$, the probability of getting the higher (or equal) prize $U_B(\tau'_S)$ is monotonically increasing in τ'_S . Therefore, the lottery with $\tau'_S = \bar{\tau}_S$ first order stochastically dominates all other lotteries in this class. The second class of lotteries is with $\tau'_S \geq \bar{\tau}_S$ and has $U_B(\bar{\tau}_B) \geq U_B(\tau'_S)$. As long as $\tau'_S \geq \bar{\tau}_S$, the probability of getting higher (or equal) prize $U_B(\bar{\tau}_B)$ is monotonically decreasing in τ'_S . Therefore, the lottery with $\tau'_S = \bar{\tau}_S$ first order stochastically dominates all other lotteries in this class. QED.

Appendix B: Example of chat and classification

| Subject and Message | Related | Code |
|---|---------|-----------|
| Subject 2: <i>So?</i> | 0 | -- |
| Subject 5: <i>What do you propose?</i> | 1 | None |
| Subject 3: <i>I propose to choose a high tax since we are the buyers.</i> | 1 | Incorrect |
| Subject 1: <i>Well, it depends on how the price is determined.</i> | 1 | None |
| Subject 2: <i>I suppose the tax will add to the price.</i> | 1 | Correct |
| Subject 5: <i>In any case, we shall all choose the same tax.</i> | 1 | None |
| Subject 1: <i>If we pay the tax anyway, why then not chose a low one?</i> | 1 | Correct |
| Subject 4: <i>All below 25?</i> | 1 | Correct |
| Subject 2: <i>I am for it.</i> | 1 | None |
| Subject 3: <i>Me too.</i> | 1 | None |
| | | |

Appendix C: Instructions for the auction (translated from German)

General Instructions for Participants

You are now participating in an economics experiment which is funded by the Austrian Science Fund. The purpose of the experiment is to analyze decision making in markets. You are paid Euro 4 for showing up on time. If you carefully read the instructions and follow the rules you can earn additional money. The €4 and all additional amounts earned during the experiment will be paid to you in cash immediately after the experiment. You can earn points in the experiment. These points will be converted to Euros according to the following exchange rate: Points 100 = €0.8. During the experiment we ask you not to speak to other participants. If you have a question, please ask us. We will gladly answer your questions in private. It is very important that you follow this rule. Otherwise the results of the experiment have no value from a scientific perspective.

You are now participating in a market experiment. In the market, you can buy units of a hypothetical commodity. You earn money by trading. How much you earn depends on your and the decisions of others. The experiment consists of two practice periods followed by a number of trading periods. In the practice periods you do not earn money; but you should take these periods seriously since you will gain valuable experience for the paid trading periods.

Detailed Instructions for Buyers

In this experiment each participant is a buyer. You can buy units from automated sellers. The sellers will sell to you according to the rules described below. In your market there are 5 buyers who can buy units from sellers in each of the trading periods.

What participants can do: As a buyer you state at which price you would buy a unit. We call this your “bid”. You can buy two units at most. You can submit separately a bid for each of these two units.

How the market works: At the end of each trading period, the bids you and other buyers in your market have submitted are collected and ranked from high to low. The highest bid is ranked above the 2nd highest bid. The 2nd highest bid is ranked above the 3rd highest bid, and so on. If two or more bids are equal, ranks will be randomly assigned by the computer.

Sellers bear costs for each unit they sell. These costs are ranked from low to high. This means that the lowest cost is ranked above the 2nd lowest cost. The 2nd lowest cost is ranked above the 3rd lowest cost, and so on. Buyers buy units with bids above costs in this order. Bids that are below costs for a unit are rejected.

Example: Assume we collect four bids in a market period. The highest bid is 145, the 2nd highest bid is 130, the 3rd highest bid is 110, and the 4th highest bid is 90. The lowest cost is 60, the 2nd lowest cost is 80, the 3rd lowest cost is 95, and the 4th lowest cost is 105.

| | 1. | 2. | 3. | 4. |
|-------|-----|-----|------|-----|
| Bids | 145 | 130 | 110* | 90 |
| Costs | 60 | 80 | 95 | 105 |

In the example the first 3 bids are above the costs of the sellers. Therefore, the 3 units sell to the buyers who have made these bids. The 4th bid at 90 is rejected. The buyer who submitted this bid does not buy the unit.

Please note: Buyers do not know the bids of others, nor do they know the costs of units for sellers.

How your earnings are computed: Each unit has a value for you. We call this “your value”. You will learn your value in the experiment. You only know your value; you do not know the values of other buyers nor do other buyers know your value. If you buy a unit you have to pay a price. Each buyer pays the same price per unit. This uniform market price is equal to the last accepted bid in the order explained above. In our example, 3 units have been bought. The bid of the last accepted unit is 110 (market with *). The uniform market price is therefore 110. Please note once again that the first two units of the order are also bought at a price of 110.

Your earnings per unit bought are calculated as follows: Earning = value minus price. Note that if you buy a unit you will pay less than your bid unless your bid is exactly equal to the market price. In our example, suppose that you submitted the bid of 130 for a unit which has a value of 155. This bid is above the market price of 110. Since you buy this unit at a market price of 110 your earnings will be $155 - 110 = 45$.

If you do not buy a unit, its value expires. On the other hand, you also do not pay a price. Your earnings are therefore zero.

Subjects’ original instructions contained figures showing how the computer screen would look like during the experiment

How is trading presented on the computer screen? In each trading period a Decision Screen appears (Figure 1). At the end of each period an Outcome Screen appears (Figure 2). After 15 trading periods a History of Results appears (Figure 3). The numbers in the figures of the instructions serve illustrative purposes only. Actual numbers may be different.

In the top area of the Decision Screen you see on the left the number of the current trading period (here: 2) and the total amount of trading periods (here: 15). Each trading period ends after a time limit. The remaining time within a period is shown in the top area on the right (here: 19 Seconds). When the experiment starts the time available for trading is generous but it will be continuously reduced in later periods.

In the first row you see your value for your first unit. In this example the buyer has a value of 155. The input field below your value serves to enter your bid. To enter a bid, click on the field labeled ‘Your Bid’ and type in a number. To submit that bid, click on the ‘Submit’ button. In the second row everything is repeated for your 2nd unit. Here, the buyer values the 2nd unit at 105 points.

Rules for bidding: An important rule for trading is to “trade at no loss”. Therefore, you may not submit bids above your values. In the example shown in Figure 1, this buyer’s bid must not be above 155 for the 1st and

105 for the 2nd unit. A second rule is that your bid for the 1st unit may not be above the bid of the 2nd unit. If you violate these rules, a message box appears. The message disappears if you press the 'OK' button.

The Outcome Screen (Figure 2) appears at the end of a trading period. The upper part of this screen looks the same as the Decision Screen. In the table below you find your value, your bid, the price, for each of the two units. In the last row of the table shows your earnings for the current period. In this example, the earnings from buying the 1st unit are 45. This is the difference between the value of this unit and the market price. The earnings from the 2nd unit are zero: the bid of 90 has been rejected because it was below the cost of the unit.

Below the table you see the total number of units traded in the market (Market Quantity). The next line shows how many units you have bought.

Finally you see Your Period Earnings. This is the sum of your earnings from buying units in the market.

After a total of 10 trading periods the History of Results (Figure 3) appears. Here, the results from each period are summarized. "Your Total Earnings over the last 10 periods" shows the sum of your earnings as sum over the last 10 periods.

Appendix D. Instructions on the tax proposal in treatment BuyerTax [*SellerTax in brackets*]

You and the 4 other buyers in your market now have to decide between two parties A and B. A tax at a rate of at least 25 points must be imposed.

Party A proposes that buyers in each of the following 10 periods pay a tax equal to 25 ($t_A=25$) [tax greater than 25 ($t_A>25$)] points for each unit they buy .

Party B proposes that sellers in each of the following 10 periods pay a tax greater than 25 ($t_B>25$) [tax equal to 25 ($t_B=25$)] points for each unit they sell.

The tax rate t_B on sellers [t_A on buyers]

At the time of your choice the tax rate t_B proposed by party B [t_A proposed by party A] is yet not known. This tax will be announced only after you have decided between the two parties A and B. The tax per unit sold [*bought*] is an integer number randomly determined between 25 and 50. All numbers are equally likely.

Decision rule

You and the 4 other buyers in your market propose a tax rate per unit. The tax rates proposed by all 5 buyers are ranked from low to high. The 3rd number in this order is called the median.

- If the median is smaller than the tax rate proposed by party B ($\text{median} < t_B$), party A wins and the tax is $t_A=25$ points: Buyers then pay a tax of 25 points on every unit they buy for the next 10 periods. [*If the median is smaller than the tax rate proposed by party A ($\text{median} < t_A$), party B wins and the tax is $t_B=25$ points: Sellers then pay a tax of 25 points on every unit they sell for the next 10 periods.*]
- If the median is greater or equal than the tax rate proposed by party B ($\text{median} \geq t_B$), party B wins and the tax is t_B points: Sellers then pay a tax of t_B points on every unit they sell for the next 10 periods. [*If the median is greater or equal than the tax rate proposed by party A ($\text{median} \geq t_A$), party A wins and the tax is t_A points: Buyers then pay a tax of t_A points on every unit they buy for the next 10 periods.*]

Please note: If you propose a tax smaller or equal than 25 you decide in favor of party A (= against party B) [*in favor of party B (= against party A)*]. The higher the median of the values proposed by you and the other 4 buyers, the more probable it is that party B [*party A*] wins.

Example 1: You and the other 4 buyers in your market propose the tax rates 35, 22, 39, 26, 37. After respective ranking, we get:

- 1: 22
- 2: 26
- 3: 35 (= median)
- 4: 37
- 5: 39

Suppose that the random tax of party B [*party A*] is determined as $t_B=32$ [$t_A=32$]. The median (35) is above $t_B=32$ [$t_A=32$]. Therefore, party B wins and sellers pay a tax of $t_B=32$ points for each unit they sell [*party A wins and buyers pay a tax of $t_A=32$ points for each unit they sell*].

Example 2: You and the other 4 buyers in your market propose the tax rates 35, 3, 17, 30, 37. After respective ranking, we get:

- 1: 3
- 2: 17
- 3: 30 (= median)
- 4: 35
- 5: 37

Suppose that the random tax of party B [*party A*] is determined at $t_B=38$ [$t_A=38$]. The median (30) is below $t_B=38$ [$t_A=38$]. Therefore, party A wins and buyers pay a tax of $t_A=25$ points for each unit they buy [*party B wins and sellers pay a tax of $t_B=25$ points for each unit they sell*].

In treatment Deliberation the following paragraph was added.

Discussion via chat

You now have the opportunity to discuss the above proposal with the other 4 buyers in your market. Discussion will be exclusively in written form via a chat on your computer. You have 5 minutes to discuss. Please, do not discuss anything apart from the proposal. For example, you may try to convince other buyers in your market about the tax rate they should choose. Please do not reveal your true identity and do not threaten or insult others. Otherwise the results of this experiment are of no value from a scientific perspective.

In treatments Info and NoInfo the following was added

Calculation of market outcomes for the next 10 periods

A large number of people have recently completed the experiment you are now participating in. We will use the data from these previous experiments to calculate market outcomes given voting choices. These market outcomes are used in turn to calculate your earnings for the next phase of the experiment. This means that you do not have to trade on the market to earn incomes in the coming phase. Instead, we calculate the income you would earn over the next 10 periods if the outcome of the vote was implemented in the market. The calculated income will be converted into money at the same exchange rate as before (100 points = 1 Euro) and paid out at the end of the experiment in cash.

Note: Extensive tests have shown that the calculated income for all possible voting outcomes is almost exactly the same the income you would have earned by trading for 10 periods. We therefore ask you to make your voting decision in the exact same way as if you were actively trading on the market for the next 10 periods.

We will proceed with the vote in a minute. Please raise your hand if you have any question.

Appendix E: Hypothetical tax scenarios (in treatments Info and NoInfo)

| | Option A | Option B |
|-------------------------------|--|---|
| Tax on beer I | Consumers have to pay an additional tax of 67.2 Cents for every liter of beer they buy. | Sellers have to pay an additional tax of 73.9 Cents for every liter of beer they sell. |
| Tax on beer II | Consumers have to pay an additional tax of 100.8 Cents for every liter of beer they buy. | Sellers have to pay an additional tax of 94.1 Cents for every liter of beer they sell. |
| Tax on insurance contracts I | Insurance holders have to pay an additional tax of 1.0 percent on the insurance premium. | Insurance companies have to pay an additional tax of 1.15 percent on the insurance premium. |
| Tax on insurance contracts II | Insurance holders have to pay an additional tax of 0.8 percent on the insurance premium. | Insurance companies have to pay an additional tax of 0.65 percent on the insurance premium. |
| Tax on electricity I | Consumers have to pay an additional tax of 1.88 Cents for every kilowatt hour of electricity they use. | Electricity companies have to pay an additional tax of 2.15 Cents for every kilowatt hour of electricity they sell. |
| Tax on electricity II | Consumers have to pay an additional tax of 1.08 Cents for every kilowatt hour of electricity they use. | Electricity companies have to pay an additional tax of 0.81 Cents for every kilowatt hour of electricity they sell. |

Notes: In treatments Info and NoInfo participants were asked to state tax preferences in 6 scenarios in phase 1 (i.e. before the first experimental vote) and phase 5 (i.e. after the second vote). The order of the scenarios, the order of the tax regimes within a scenario, and whether Option A or B had the higher tax was randomized.

Appendix F: Exogenously provided information on consequences of tax regimes in phase 3 of treatment Info

| Median | Seller tax | Implemented tax s: seller tax regime b: buyer tax regime | Market price | Market income (10 periods) |
|--------|------------|--|--------------|-------------------------------|
| 50 | 46 | s = 46 | 138 | 171 |
| 4 | 27 | b = 25 | 90 | 398 |
| 30 | 30 | s = 30 | 122 | 335 |
| 49 | 29 | s = 29 | 121 | 344 |
| 23 | 28 | b = 25 | 91 | 394 |
| 39 | 44 | b = 25 | 95 | 352 |
| 10 | 40 | b = 25 | 94 | 365 |
| 35 | 37 | b = 25 | 92 | 383 |
| 40 | 39 | s = 39 | 134 | 213 |

Notes: In treatment Info, information on 9 median votes from previous experiments was presented before the second vote (numbers in columns Median and Seller tax were taken from actual realizations in BuyerTax. Numbers in column Implemented tax were determined by application of the random mechanism described in section 2.2. Numbers in columns Market price and Market income are based on equilibrium predictions plus noise (uniformly distributed random variable between 0 and 5 which was added to the market price).